

Strong reductions for extended formulations

Gábor Braun¹, Sebastian Pokutta², and Aurko Roy³

¹ISyE, Georgia Institute of Technology, Atlanta, GA, USA. *Email:* gabor.braun@isye.gatech.edu

²ISyE, Georgia Institute of Technology, Atlanta, GA, USA. *Email:* sebastian.pokutta@isye.gatech.edu

³College of Computing, Georgia Institute of Technology, Atlanta, GA, USA. *Email:* aurko@gatech.edu

December 15, 2015

Abstract

We generalize the reduction mechanism for linear programming problems and semidefinite programming problems from [BPZ15] in two ways (1) relaxing the requirement of affineness, and (2) extending to fractional optimization problems.

As applications we provide several new LP-hardness and SDP-hardness results, e.g., for the SparsestCut problem, the BalancedSeparator problem, the MaxCut problem and the Matching problem on 3-regular graphs. We also provide a new, very strong Lasserre integrality gap for the IndependentSet problem, which is strictly greater than the best known LP approximation, showing that the Lasserre hierarchy does not always provide the tightest SDP relaxation.

1 Introduction

Linear and semidefinite programs are the main components in the design of many practical algorithms and therefore understanding their expressive power is a fundamental problem. The complexity of these programs is measured by the number of constraints, ignoring all other aspects affecting the running time of an actual algorithm, in particular, these measures are independent of the P vs. NP question. We call a problem *LP-hard* if it does not admit an LP formulation with a polynomial number of constraints, and we define *SDP-hardness* similarly.

Recently, motivated by Yannakakis’s influential work [Yan88, Yan91], a plethora of strong lower bounds have been established for many important optimization problems, such as e.g., the Matching problem [Rot14] or the TravelingSalesman problem [FMP⁺12, FMP⁺15, LRS14]. In [BPZ15], the authors introduced a reduction mechanism providing inapproximability results for large classes of problems. However, the reductions were required to be affine, and hence failed for e.g., the VertexCover problem, where intermediate Sherali–Adams reductions were employed in [BFPS15] due to this shortcoming.

In this work we extend the reduction mechanism of [BPZ15] in two ways, establishing several new hardness results both in the LP and SDP setting; both are special cases arising from reinterpreting LPs and SDPs as proof systems (see Section 2.2). First, by including additional ‘computation’ in the reduction, we allow non-affine relations between problems, eliminating the need for Sherali–Adams reductions in [BFPS15]. Second, we extend the framework to fractional optimization problems (such as e.g., SparsestCut) where ratios of linear functions have to be optimized. Here typically one optimizes the numerator and denominator at the same time, and that is what we incorporate in our framework.

Related Work

The immediate precursor to this work is [BPZ15] (generalizing [Pas12, BFPS12]), introducing a reduction mechanism. Base hard problems are the Matching problem [Rot14], as well as constraint satisfaction problems [CLRS13, LRS14] based on hierarchy hardness results, such as e.g., [Sch08] and [CMM09].

Contribution

Generalized LP/SDP reductions. We generalize the reduction mechanism in [BPZ15] by modeling additional computation, i.e., using extra LP or SDP constraints. Put differently, we allow for more complicated reduction maps as long as these maps themselves have a small LP/SDP formulation. As a consequence, we can relax the affineness requirement and enable a weak form of *gap-amplification* and *boosting*. This overcomes a major limitation of the approach in [BPZ15], yielding significantly stronger reductions at a small cost.

Fractional LP/SDP optimization. Second, we present a *fractional* LP/SDP framework and reduction mechanism, where the objective functions are ratios of functions from a low dimensional space, such as for the SparsestCut problem. For these problems the standard LP/SDP framework is meaningless as the ratios span a high dimensional affine space. The fractional framework models the usual way of solving fractional optimization problems, enabling us to establish strong statements about LP or SDP complexity.

Direct non-linear hardness reductions. We demonstrate the power of our generalized reduction by establishing new LP-hardness and SDP-hardness for several problems of interest, i.e., these problems cannot be solved by LPs/SDPs of polynomial size; see Table 1. We establish various hardness results for the SparsestCut and BalancedSeparator problems even when one of the underlying graph has bounded treewidth. We redo the reductions to intermediate CSP problems used for optimal inapproximability results for the VertexCover problem over simple graphs and Q -regular hypergraphs in [BFPS15], eliminating Sherali–Adams reductions. We also show the first explicit SDP-hardness for the MaxCut problem, inapproximability within a factor of $15/16 + \epsilon$, which is stronger than the algorithmic hardness of $16/17 + \epsilon$. Finally, we prove a new, strong Lasserre integrality gap of $n^{1-\gamma}$ after $O(n^\gamma)$ rounds for the IndependentSet problem for any sufficiently small $\gamma > 0$. It not only significantly strengthens and complements the best-known integrality gap results so far ([Tul09] and [AT11, AT13]; see also [LT03, ST99]), but also shows the suboptimality of Lasserre relaxations for the IndependentSet problem together with [BFPS15].

Small uniform LPs for bounded treewidth problems. Finally, we introduce a new technique in Section 10 to derive small *uniform* linear programs for problems over graphs of bounded treewidth. Here the same linear program is used for all bounded treewidth instances of the same size, independent of the actual tree decompositions, whereas the linear program in [KKT15] work for a single input instance only (with fewer inequalities than our linear program).

Outline

We start by recalling and refining the linear programming framework in Section 2, including the optimization problems we shall consider. We develop a general theory in Section 2.2 leading easily to both a generalized reduction mechanism in Section 3 and an extension to fractional optimization in Section 4. The remaining chapters contain mostly applications to various problems. Exceptions are Section 8, establishing a Lasserre integrality gap for the IndependentSet problem, and Section 10 providing a small linear program for bounded treewidth problems.

Problem	Factor	Source	Paradigm	Remark
MaxCut	$\frac{15}{16} + \varepsilon$	Max-3-XOR/0	SDP	
SparsestCut(n), $\text{tw}(\text{supply}) = O(1)$	$2 - \varepsilon$	MaxCut	LP	opt. [CLRS13]
SparsestCut(n), $\text{tw}(\text{supply}) = O(1)$	$\frac{16}{15} - \varepsilon$	MaxCut	SDP	
BalancedSeparator(n, d), $\text{tw}(\text{demand}) = O(1)$	$\omega(1)$	UniqueGames	LP	
IndependentSet	$\omega(n^{1-\varepsilon})$	Max- k -CSP	Lasserre $O(n^\varepsilon)$ rounds	
Matching, 3-regular	$1 + \varepsilon/n^2$	Matching	LP	
1F-CSP	$\omega(1)$	UniqueGames	LP	[BFPS15] w/o SA
$Q \neq$ -CSP				

Table 1: Inapproximability of optimization problems. tw denotes treewidth.

2 Preliminaries

Here we recall the linear programming and semidefinite programming framework from [BPZ15], as well as the optimization problems we shall consider later, paying particular attention to base hard problems. Section 2.2 is a new technical foundation for the framework, presenting the underlying theory in a unified simple way, from which the extensions in Sections 3 and 4 readily follow. We start by recalling the notion of *tree decompositions* and *treewidth* of a graph.

Definition 2.1 (Tree width). A *tree decomposition* of a graph G is a tree T together with a vertex set of G called *bag* $B_t \subseteq V(G)$ for every node t of T , satisfying the following conditions: (1) $V(G) = \bigcup_{t \in V(T)} B_t$, (2) For every adjacent vertices u, v of G there is a bag B_t containing both u and v , and (3) For all nodes t_1, t_2, t of T with t lying between t_1 and t_2 (i.e., t is on the unique path connecting t_1 and t_2) we have $B_{t_1} \cap B_{t_2} \subseteq B_t$. The *width* of the tree decomposition is $\max_{t \in V(T)} |B_t| - 1$: one less than the maximum bag size. The *treewidth* $\text{tw}(G)$ of G is the minimum width of its tree decompositions.

We will use $\chi(\cdot)$ for indicator functions: i.e., $\chi(X) = 1$ if the statement X is true, and $\chi(X) = 0$ otherwise. We will denote random variables using bold face, e.g. \mathbf{x} . Let S^r denote the set of symmetric $r \times r$ real matrices, and let S^r_+ denote the set of positive semidefinite $r \times r$ real matrices.

2.1 Optimization Problems

Definition 2.2 (Optimization problem). An *optimization problem* is a tuple $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ consisting of a set \mathcal{S} of *feasible solutions*, a set \mathcal{I} of *instances*, and a real-valued objective called *measure* $\text{val}: \mathcal{I} \times \mathcal{S} \rightarrow \mathbb{R}$.

We shall write $\text{val}_{\mathcal{I}}(s)$ for the objective value of a feasible solution $s \in \mathcal{S}$ for an instance $\mathcal{I} \in \mathcal{I}$.

The SparsestCut problem is defined over a graph with two kinds of edges: *supply* and *demand* edges. The objective is to find a cut that minimizes the ratio of the capacity of cut supply edges to the total demand separated. For a weight function $f: E(K_n) \rightarrow \mathbb{R}_{\geq 0}$, we define the graph $[n]_f := ([n], E_f)$ where $E_f := \{(i, j) \mid i, j \in [n], f(i, j) > 0\}$. We study the SparsestCut problem with bounded-treewidth supply graph.

Definition 2.3 (SparsestCut(n, k)). Let n be a positive integer. The minimization problem SparsestCut(n, k) consists of

instances a pair (d, c) of a nonnegative *demand* $d: E(K_n) \rightarrow \mathbb{R}_{\geq 0}$ and a *capacity* $c: E(K_n) \rightarrow \mathbb{R}_{\geq 0}$ such that $\text{tw}([n]_c) \leq k$;

feasible solutions all subsets s of $[n]$;

measure ratio of separated capacity and separated demand:

$$\text{val}_{d,c}(s) = \frac{\sum_{i \in s, j \notin s} c(i, j)}{\sum_{i \in s, j \notin s} d(i, j)}$$

for capacity c , demand d , and set s .

The **BalancedSeparator** problem is similar to the **SparsestCut** problem and is also defined over a graph with *supply* and *demand* edges. However it restricts the solutions to cuts that are *balanced*, i.e., which separate a large proportion of the demand. Note that in this case we define the **BalancedSeparator** problem on n vertices for a fixed demand function d , unlike in the case of **SparsestCut** where the demand function d was part of the instances. This is because in the framework of [BPZ15] the solutions should be independent of the instances. We formalize this below.

Definition 2.4 (**BalancedSeparator**(n, d)). Let n be a positive integer, and $d: [n] \times [n] \rightarrow \mathbb{R}_{\geq 0}$ a nonnegative function called *demand function*. Let D denote the total demand $\sum_{i,j \in [n]} d(i, j)$. The minimization problem **BalancedSeparator**(n, d) consists of

instances nonnegative *capacity* function $c: E(K_n) \rightarrow \mathbb{R}_{\geq 0}$ on the edges of the complete graph K_n ;

feasible solutions all subsets s of $[n]$ such that $\sum_{i \in s, j \notin s} d(i, j)$ is at least $D/4$;

measure capacity of cut supply edges: $\text{val}_c(s) := \sum_{i \in s, j \notin s} c(i, j)$ for a capacity function c and set s .

Recall that an *independent set* I of a graph G is a subset of pairwise non-adjacent vertices $I \subseteq V(G)$. The **IndependentSet** problem on a graph G asks for an independent set of G of maximum size. We formally define it as an optimization problem below.

Definition 2.5 (**IndependentSet**(G)). Given a graph G , the maximization problem **IndependentSet**(G) consists of

instances all induced subgraphs H of G ;

feasible solutions all independent subsets I of G ;

measure $\text{val}_H(I) = |I \cap V(H)|$.

Recall that a subset X of $V(G)$ for a graph G is a *vertex cover* if every edge of G has at least one end point in X . The **VertexCover** problem on a graph G asks for a vertex cover of G of minimum size. We give a formal definition below.

Definition 2.6 (**VertexCover**(G)). Given a graph G , the minimization problem **VertexCover**(G) consists of

instances all induced subgraphs H of G ;

feasible solutions all vertex covers X of G ;

measure $\text{val}_H(X) = |X \cap V(H)|$.

The **MaxCut** problem on a graph G asks for a vertex set of G cutting a maximum number of edges. Given a vertex set $X \subseteq V(G)$, let $\delta_G(X) := \{\{u, v\} \in E(G) \mid u \in X, v \notin X\}$ denote the set of edges of G with one end point in X and the other end point outside X .

Definition 2.7 (**MaxCut**(G)). Given a graph G , the maximization problem **MaxCut**(G) consists of

instances all induced subgraph H of G ;

feasible solutions all vertex subsets $X \subseteq V(G)$;

measure $\text{val}_H(X) = |E(H) \cap \delta_G(X)|$.

Constraint satisfaction problems (CSPs for short) are inherently related to inapproximability results, and form a basic collection of inapproximable problems. There are many variants of CSPs, but the general structure is as follows:

Definition 2.8 (Constraint Satisfaction Problems). A *constraint satisfaction problem*, in short CSP, is an optimization problem on a fixed set $\{x_1, \dots, x_n\}$ of variables with values in a fixed set $[q]$ consisting of

instances formal weighted sums $\mathcal{I} = \sum_i w_i C_i(x_{j_1}, \dots, x_{j_{k_i}})$ of some clauses $C_i: [q]^{k_i} \rightarrow \{0, 1\}$ with weights $w_i \geq 0$.

feasible solutions all mappings $s: \{x_1, \dots, x_n\} \rightarrow [q]$, called *assignments* to variables

measure weighted fraction of satisfied clauses:

$$\text{val}_{\mathcal{I}}(s) := \frac{\sum_i w_i C_i(s(x_{j_1}), \dots, s(x_{j_{k_i}}))}{\sum_i w_i}$$

A CSP can be either a maximization problem or a minimization problem. For specific CSPs there are restrictions on permitted clauses, and later we will define CSPs by specifying only these restrictions. For example Max- k -CSP is the problem where only clauses with at most k free variables are allowed (i.e., $k_i \leq k$ in the definition above). The problem Max- k -XOR is the problem with clauses of the form $x_1 + \dots + x_k = b$ where the x_i are distinct variables, $b \in \{0, 1\}$, and the addition is modulo 2. We shall use the subproblem Max- k -XOR/0, where the clauses have the form $x_1 + \dots + x_k = 0$.

Given a k -ary predicate P , let Max- k -CSP(P) denote the CSP where all clauses arise via a change of variables from P , i.e., every clause have the form $P(x_{i_1}, \dots, x_{i_k})$ with i_1, \dots, i_k being pairwise distinct. For example, Max- k -XOR/0 = Max- k -CSP($x_1 + \dots + x_k = 0$).

Another specific example of a CSP we will make use of is the UniqueGames problem. The UniqueGames problem asks for a labeling of the vertices of a graph that maximizes the number (or weighted sum) of edges where the labels of the endpoints match. We formalize it restricted to regular bipartite graphs.

Definition 2.9 (UniqueGames $_{\Delta}(n, q)$). Let n, q and Δ be positive integer parameters. The maximization problem UniqueGames $_{\Delta}(n, q)$ consists of

instances All edge-weighted Δ -regular bipartite graphs (G, w) (i.e., a graph G with a collection $\{w_{u,v}\}_{\{u,v\} \in E(G)}$ of real numbers) with partite sets $\{0\} \times [n]$ and $\{1\} \times [n]$ with every edge $\{i, j\}$ labeled with a permutation $\pi_{i,j}: [q] \rightarrow [q]$ such that $\pi_{i,j} = \pi_{j,i}^{-1}$.

feasible solutions All functions $s: \{0, 1\} \times [n] \rightarrow [q]$ called *labelings* of the vertices.

measure The weighted fraction of correctly labeled edges, i.e., edges $\{i, j\}$ with $s(i) = \pi_{i,j}(s(j))$:

$$\text{val}_{(G,w)}(s) := \frac{\sum_{\substack{\{i,j\} \in E(G) \\ s(i) = \pi_{i,j}(s(j))}} w(i, j)}{\sum_{\{i,j\} \in E(G)} w(i, j)}$$

The Matching problem asks for a matching in a graph H of maximal size. The restriction to matchings and subgraphs (which corresponds to 0/1 weights in the objective of the matching problem) below serves the purpose to obtain a base hard problem, with which we can work more easily later.

Definition 2.10 (Matching(G)). The maximum matching problem Matching(G) over a graph G is defined as the maximization problem:

instances all subgraphs H of G

feasible solutions all perfect matchings S on G .

measure the size of induced matching $\text{val}_G(S) := |S \cap E(H)|$ with $S \in \mathcal{S}$, and H a subgraph of G .

We will also write Matching $_k(G)$ to indicate that the maximum vertex degree is at most k .

2.1.1 Uniform problems

Here we present so called *uniform* versions of some of the optimization problems discussed so far, where the class of instances is typically much larger, e.g., the class of all instances of a given size. *Non-uniform* optimization problems typically consider weighted versions of a *specific instance* or all induced subgraphs of a *given graph*. For establishing lower bounds, non-uniform optimization problems give stronger bounds: ‘even if we consider a *specific graph*, then there is no small LP/SDP’. In the case of upper bounds, i.e., when we provide formulations, uniform optimization problems provide stronger statements: ‘even if we consider *all graphs simultaneously*, then there exists a small LP/SDP’.

We will later show in Section 10 that over graphs of bounded tree-width there exists a small LP that solves the uniform version of optimization problems. We start by defining the uniform version of MaxCut. Recall that for a graph G and a subset X of $V(G)$, we define $\delta_G(X) := \{\{u, v\} \in E(G) \mid u \in X, v \notin X\}$ to be the set of crossing edges.

Definition 2.11 (MaxCut(n)). For a positive integer n , the maximization problem MaxCut(n) consists of

instances all graphs G with $V(G) \subseteq [n]$;

feasible solutions all subsets X of $[n]$;

measure $\text{val}_G(X) = |\delta_G(X)|$.

With IndependentSet and VertexCover we face the difficulty that the solutions are instance dependent. Hence we enlarge the feasible solutions to include all possible vertex sets, and in the objective function penalize the violation of requirements.

Definition 2.12 (IndependentSet(n)). For a positive integer n , the maximization problem IndependentSet(n) consists of

instances all graphs G with $V(G) \subseteq [n]$;

feasible solutions all subsets X of $[n]$;

measure the number of vertices of G in X penalized by the number of edges of G inside X :

$$\text{val}_G(X) = |X \cap V(G)| - |E(G[X])|. \quad (1)$$

Recall that VertexCover asks for a minimal size vertex set X of a graph G such that every edge of G has at least one of its endpoints in X .

Definition 2.13. For a positive integer n the minimization problem VertexCover consists of

instances all graphs G with $V(G) \subseteq [n]$

feasible solutions all subsets $X \subseteq V(G)$

measure the number of vertices of G in X penalized by the number of uncovered edges:

$$\text{val}_G(X) := |X \cap V(G)| + |E(G \setminus X)|. \quad (2)$$

2.2 Nonnegativity problems: Extended formulations as proof system

In this section we introduce an abstract view of formulation complexity, where the main idea is to reduce all statements to the core question about the complexity of deriving nonnegativity for a class of nonnegative functions. This abstract view will allow us to easily introduce future versions of reductions and optimization problems with automatic availability of Yannakakis's Factorization Theorem and the reduction mechanism.

Definition 2.14. A nonnegativity problem $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ consists of a set \mathcal{I} of instances, a set \mathcal{S} of feasible solutions and a nonnegative evaluation $\text{val}: \mathcal{I} \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$.

As before, we shall write $\text{val}_{\mathcal{I}}(s)$ instead of $\text{val}(\mathcal{I}, s)$. The aim is to study the complexity of proving nonnegativity of the functions $\text{val}_{\mathcal{I}}$. Therefore we define the notion of proof as a linear program or a semidefinite program.

Definition 2.15. Let $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ be a nonnegativity problem. An LP proof of nonnegativity of \mathcal{P} consists of a linear program $Ax \leq b$ with $x \in \mathbb{R}^r$ for some r and the following realizations:

Feasible solutions as vectors $x^s \in \mathbb{R}^r$ for every $s \in \mathcal{S}$ satisfying

$$Ax^s \leq b \quad \text{for all } s \in \mathcal{S}, \quad (3)$$

i.e., the system $Ax \leq b$ is a relaxation (superset) of $\text{conv}(x^s \mid s \in \mathcal{S})$.

Instances as affine functions $w_{\mathcal{I}}: \mathbb{R}^r \rightarrow \mathbb{R}$ for all $\mathcal{I} \in \mathcal{I}^{\mathcal{S}}$ satisfying

$$w_{\mathcal{I}}(x^s) = \text{val}_{\mathcal{I}}(s) \quad \text{for all } s \in \mathcal{S}, \quad (4)$$

i.e., the linearization $w_{\mathcal{I}}$ of $\text{val}_{\mathcal{I}}$ is required to be exact on all x^s with $s \in \mathcal{S}$.

Proof We require that the $w_{\mathcal{I}}$ are nonnegative on the solution set of the LP:

$$w_{\mathcal{I}}(x) \geq 0 \quad \text{whenever } Ax \leq b, \mathcal{I} \in \mathcal{I}. \quad (5)$$

The size of the formulation is the number of inequalities in $Ax \leq b$. Finally, LP proof complexity $\text{fc}_{\text{LP}}(\mathcal{P})$ of \mathcal{P} is the minimal size of all its LP proofs.

The notion of an SDP proof is defined similarly.

Definition 2.16. Let $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ be a nonnegativity problem. An SDP proof of nonnegativity of \mathcal{P} consists of a semidefinite program $\{X \in \mathbb{S}_+^r \mid \mathcal{A}(X) = b\}$ (i.e., a linear map $\mathcal{A}: \mathbb{S}^r \rightarrow \mathbb{R}^k$ together with a vector $b \in \mathbb{R}^k$) and the following realizations:

Feasible solutions as vectors $X^s \in \mathbb{S}_+^r$ for all $s \in \mathcal{S}$ satisfying

$$\mathcal{A}(X^s) = b \quad (6)$$

Instances as nonnegative affine functions $w_{\mathcal{I}}: \mathbb{S}^r \rightarrow \mathbb{R}$ for all $\mathcal{I} \in \mathfrak{I}$ satisfying

$$w_{\mathcal{I}}(X^s) = \text{val}_{\mathcal{I}}(s) \quad \text{for all } s \in \mathcal{S}. \quad (7)$$

Proof We require nonnegativity on the feasible region of the SDP:

$$w_{\mathcal{I}}(X) \geq 0 \quad \text{whenever } \mathcal{A}(X) = b, X \in \mathbb{S}_+^r, \mathcal{I} \in \mathfrak{I}. \quad (8)$$

The *size* of the formulation is the dimension parameter r . Finally, the *SDP proof complexity* $\text{fc}_{\text{SDP}}(\mathcal{P})$ of \mathcal{P} is the minimal size of all its SDP proofs.

2.2.1 Slack matrix and proof complexity

We introduce the slack matrix of a nonnegativity problem as a main tool to study proof complexity, generalizing the approach from the polyhedral world. The main result is a version of Yannakakis's Factorization Theorem formulating proof complexity in the language of linear algebra as a combinatorial property of the slack matrix.

Definition 2.17. The *slack matrix* of a nonnegativity problem $\mathcal{P} = (\mathcal{S}, \mathfrak{I}, \text{val})$ is the $\mathfrak{I} \times \mathcal{S}$ matrix $M_{\mathcal{P}}$ with entries the values of the function $\text{val}_{\mathcal{I}}$

$$M_{\mathcal{P}}(\mathcal{I}, s) := \text{val}_{\mathcal{I}}(s). \quad (9)$$

We will use the standard notions of nonnegative rank and semidefinite rank.

Definition 2.18 ([BPZ15]). Let M be a nonnegative matrix.

nonnegative factorization A *nonnegative factorization* of M of size r is a decomposition $M = \sum_{i=1}^r M_i$ of M as a sum of r nonnegative matrices M_i of rank 1. The *nonnegative rank* $\text{rk}_+ M$ is the minimum r for which M has a nonnegative factorization of size r .

psd factorization A *positive semi-definite (psd) factorization* of M of size r is a decomposition $M(\mathcal{I}, s) = \text{Tr}[A_{\mathcal{I}} B_s]$ of M where the $A_{\mathcal{I}}$ and B_s are positive semi-definite (psd) $r \times r$ matrices. The *psd rank* $\text{rk}_{\text{psd}} M$ is the minimum r for which M has a psd factorization of size r .

We define variants ignoring factors of the form $a\mathbb{1}$:

LP factorization An *LP factorization* of M of size r is a decomposition $M = \sum_{i=1}^r M_i + u\mathbb{1}$ of M as a sum of r nonnegative matrices M_i of rank 1 and possibly an additional nonnegative rank-1 $u\mathbb{1}$ with all columns being equal. The *LP rank* $\text{rk}_{\text{LP}} M$ is the minimum r for which M has an LP factorization of size r .

SDP factorization An *SDP factorization* of M of size r is a decomposition $M(\mathcal{I}, s) = \text{Tr}[A_{\mathcal{I}} B_s] + u_{\mathcal{I}}$ of M where the $A_{\mathcal{I}}$ and B_s are positive semi-definite (psd) $r \times r$ matrices, and $u_{\mathcal{I}}$ is a nonnegative number. The *SDP rank* $\text{rk}_{\text{SDP}} M$ is the minimum r for which M has an SDP factorization of size r .

Remark 2.19. The difference between LP rank and nonnegative rank (see Definition 2.18) is solely by measuring the size of a factorization: for LP rank factors with equal columns do not contribute to the size. This causes a difference of at most 1 between the two ranks. The motivation for the LP rank is that it captures exactly the LP formulation complexity of an optimization problem, in particular for approximation problems (see [BPZ15] for an in-depth discussion). Similar remarks apply to the relation of SDP rank, psd rank, and SDP formulation complexity.

Theorem 2.20. *For every nonnegativity problem \mathcal{P} with slack matrix $M_{\mathcal{P}}$ we have*

$$\text{fc}_{\text{LP}}(\mathcal{P}) = \text{rk}_{\text{LP}} M_{\mathcal{P}}, \quad (10)$$

$$\text{fc}_{\text{SDP}}(\mathcal{P}) = \text{rk}_{\text{SDP}} M_{\mathcal{P}}, \quad (11)$$

Proof. The proof is an extension of the usual proofs of Yannakakis's Factorization Theorem, e.g., that in [BPZ15]. We provide the proof only for the LP case, as the proof for the SDP case is similar.

First we prove $\text{rk}_{\text{LP}} M_{\mathcal{P}} \leq \text{fc}_{\text{LP}}(\mathcal{P})$. Let $Ax \leq b$ be an LP proof for \mathcal{P} of size $\text{fc}_{\text{LP}}(\mathcal{P})$ with realization x^s for $s \in \mathcal{S}$ and affine functions $w_{\mathcal{I}}$ for $\mathcal{I} \in \mathcal{I}$. By Farkas's lemma, there are nonnegative matrices $u_{\mathcal{I}}$ and nonnegative numbers $\gamma_{\mathcal{I}}$ with $w_{\mathcal{I}}(x) = u_{\mathcal{I}} \cdot (b - Ax) + \gamma_{\mathcal{I}}$. Substituting x by x^s , we obtain an LP factorization of size $\text{fc}_{\text{LP}}(\mathcal{P})$:

$$M_{\mathcal{P}}(\mathcal{I}, s) = \text{val}_{\mathcal{I}}(s) = w_{\mathcal{I}}(x^s) = u_{\mathcal{I}} \cdot (b - Ax^s) + \gamma_{\mathcal{I}}.$$

Conversely, to show $\text{fc}_{\text{LP}}(\mathcal{P}) \leq \text{rk}_{\text{LP}}(\mathcal{P})$, we choose an LP factorization of $M_{\mathcal{P}}$ of size $r = \text{rk}_{\text{LP}}(\mathcal{P})$

$$M_{\mathcal{P}}(\mathcal{I}, s) = u_{\mathcal{I}} x^s + \gamma_{\mathcal{I}}$$

where the $u_{\mathcal{I}}$ and x^s are nonnegative matrices of size $1 \times r$ and $r \times 1$, respectively, and the $\gamma_{\mathcal{I}}$ are nonnegative numbers. Now \mathcal{P} has the following LP proof: The linear program is $x \geq 0$ for $x \in \mathbb{R}^{r \times 1}$. A feasible solution s is represented by the vector x^s . An instance \mathcal{I} is represented by

$$w_{\mathcal{I}}(x) := u_{\mathcal{I}} x + \gamma_{\mathcal{I}}.$$

To check the proof, note that by nonnegativity of $u_{\mathcal{I}}$ and $\gamma_{\mathcal{I}}$, we have $w_{\mathcal{I}}(x) \geq 0$ for all $x \geq 0$. Clearly, $w_{\mathcal{I}}(x^s) = M_{\mathcal{P}}(\mathcal{I}, s) = \text{val}_{\mathcal{I}}(s)$, completing the proof. \square

2.2.2 Reduction between nonnegativity problems

Definition 2.21 (Reduction). Let $\mathcal{P}_1 = (\mathcal{S}_1, \mathcal{I}_1, \text{val}^{\mathcal{P}_1})$ and $\mathcal{P}_2 = (\mathcal{S}_2, \mathcal{I}_2, \text{val}^{\mathcal{P}_2})$ be nonnegativity problems.

A reduction from \mathcal{P}_1 to \mathcal{P}_2 consists of

1. two mappings: $*$: $\mathcal{I}_1 \rightarrow \mathcal{I}_2$ and $*$: $\mathcal{S}_1 \rightarrow \mathcal{S}_2$ translating instances and feasible solutions independently;
2. two nonnegative $\mathcal{I}_1 \times \mathcal{S}_1$ matrices M_1, M_2

satisfying

$$\text{val}_{\mathcal{I}_1}^{\mathcal{P}_1}(s_1) = \text{val}_{\mathcal{I}_2}^{\mathcal{P}_2}(s_1^*) \cdot M_1(\mathcal{I}_1, s_1) + M_2(\mathcal{I}_1, s_1). \quad (12)$$

The matrices M_1 and M_2 encode additional arguments in the nonnegativity proof of \mathcal{P}_1 , besides using nonnegativity of \mathcal{P}_2 . Therefore in applications they should have low complexity, to provide a strong reduction. The following theorem relates the proof complexity of problems in a reduction.

Theorem 2.22. *Let \mathcal{P}_1 and \mathcal{P}_2 be nonnegativity problems with a reduction from \mathcal{P}_1 to \mathcal{P}_2 . Then*

$$\text{fc}_{\text{LP}}(\mathcal{P}_1) \leq \text{rk}_{\text{LP}} M_2 + \text{rk}_{\text{LP}} M_1 + \text{rk}_+ M_1 \cdot \text{fc}_{\text{LP}}(\mathcal{P}_2), \quad (13)$$

$$\text{fc}_{\text{SDP}}(\mathcal{P}_1) \leq \text{rk}_{\text{SDP}} M_2 + \text{rk}_{\text{SDP}} M_1 + \text{rk}_{\text{psd}} M_1 \cdot \text{fc}_{\text{SDP}}(\mathcal{P}_2), \quad (14)$$

where M_1 and M_2 are the matrices in the reduction as in Definition 2.21.

Proof. We prove the claim only for the LP rank, as the proof for the SDP rank is similar. We apply the Factorization Theorem (Theorem 2.20). Let $M_{\mathcal{P}_1}$ and $M_{\mathcal{P}_2}$ denote the slack matrices of \mathcal{P}_1 and \mathcal{P}_2 , respectively. Then Eq. (12) can be written as

$$M_{\mathcal{P}_1} = (F_{\mathcal{I}} M_{\mathcal{P}_2} F_{\mathcal{S}}) \circ M_1 + M_2, \quad (15)$$

where \circ denotes the Hadamard product (entrywise product), and $F_{\mathcal{I}}$ and $F_{\mathcal{S}}$ are the $\mathcal{I}_1 \times \mathcal{I}_2$ and $\mathcal{S}_2 \times \mathcal{S}_1$ matrices encoding the two maps $*$, respectively:

$$F_{\mathcal{I}}(\mathcal{I}_1, \mathcal{I}_2) := \begin{cases} 1 & \text{if } \mathcal{I}_2 = \mathcal{I}_1^*, \\ 0 & \text{if } \mathcal{I}_2 \neq \mathcal{I}_1^*; \end{cases} \quad F_{\mathcal{S}}(\mathcal{S}_2, \mathcal{S}_1) := \begin{cases} 1 & \text{if } \mathcal{S}_2 = \mathcal{S}_1^*, \\ 0 & \text{if } \mathcal{S}_2 \neq \mathcal{S}_1^*. \end{cases} \quad (16)$$

Let $M_{\mathcal{P}_2} = \tilde{M}_{\mathcal{P}_2} + a\mathbb{1}$ with $\text{rk}_{\text{LP}} M_{\mathcal{P}_2} = \text{rk}_+ \tilde{M}_{\mathcal{P}_2}$. This enables us to further simplify Eq. (15):

$$M_{\mathcal{P}_1} = (F_{\mathcal{I}} \tilde{M}_{\mathcal{P}_2} F_{\mathcal{S}}) \circ M_1 + \text{diag}(F_{\mathcal{I}} a) \cdot M_1 + M_2, \quad (17)$$

where $\text{diag}(x)$ stands for the square diagonal matrix with the entries of x in the diagonal. Now the claim follows from Theorem 2.20, the well-known identities $\text{rk}_+(A \circ B) \leq \text{rk}_+ A \cdot \text{rk}_+ B$, $\text{rk}_+ ABC \leq \text{rk}_+ B$, and the obvious $\text{rk}_{\text{LP}}(A + B) \leq \text{rk}_{\text{LP}} A + \text{rk}_{\text{LP}} B$ together with $\text{rk}_{\text{LP}}(AB) \leq \text{rk}_{\text{LP}} B$. \square

2.3 LP and SDP formulations

Here we recall the notion of linear programming and semi-definite programming complexity of optimization problems from [BPZ15]. The key idea to modeling approximations of an optimization problem $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ is to represent the approximation gap by two functions $C, S: \mathcal{I} \rightarrow \mathbb{R}$, the *completeness guarantee* and *soundness guarantee*, respectively, and the task is to differentiate problems with $\text{OPT}(\mathcal{I}) \leq S(\mathcal{I})$ and $\text{OPT}(\mathcal{I}) \geq C(\mathcal{I})$, as in the algorithmic setting.

The guarantees C and S will often be of the form $C = \alpha g$ and $S = \beta g$ for some constants α and β and an easy-to-compute function g . Then we shall write $\text{f}_{\text{LP}}(\mathcal{P}, \alpha, \beta)$ instead of the more precise $\text{f}_{\text{LP}}(\mathcal{P}, \alpha g, \beta g)$.

Definition 2.23 (LP formulation of an optimization problem). Let $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ be an optimization problem, and C, S be real-valued functions on \mathcal{I} , called *completeness guarantee* and *soundness guarantee*, respectively. If \mathcal{P} is a maximization problem, then let $\mathcal{I}^S := \{\mathcal{I} \in \mathcal{I} \mid \max \text{val}_{\mathcal{I}} \leq S(\mathcal{I})\}$ denote the set of instances, for which the soundness guarantee S is an upper bound on the maximum. If \mathcal{P} is a minimization problem, then let $\mathcal{I}^S := \{\mathcal{I} \in \mathcal{I} \mid \min \text{val}_{\mathcal{I}} \geq S(\mathcal{I})\}$ denote the set of instances, for which the soundness guarantee S is a lower bound on the minimum.

A (C, S) -approximate LP formulation of \mathcal{P} consists of a linear program $Ax \leq b$ with $x \in \mathbb{R}^r$ for some r and the following *realizations*:

Feasible solutions as vectors $x^s \in \mathbb{R}^r$ for every $s \in \mathcal{S}$ satisfying

$$Ax^s \leq b \quad \text{for all } s \in \mathcal{S}, \quad (18)$$

i.e., the system $Ax \leq b$ is a relaxation of $\text{conv}(x^s \mid s \in \mathcal{S})$.

Instances as affine functions $w_{\mathcal{I}}: \mathbb{R}^r \rightarrow \mathbb{R}$ for all $\mathcal{I} \in \mathcal{I}^S$ satisfying

$$w_{\mathcal{I}}(x^s) = \text{val}_{\mathcal{I}}(s) \quad \text{for all } s \in \mathcal{S}, \quad (19)$$

i.e., the linearization $w_{\mathcal{I}}$ of $\text{val}_{\mathcal{I}}$ is required to be exact on all x^s with $s \in \mathcal{S}$.

Achieving (C, S) approximation guarantee by requiring

$$\max \{w_{\mathcal{I}}(x) \mid Ax \leq b\} \leq C(\mathcal{I}) \quad \text{for all } \mathcal{I} \in \mathcal{I}^S, \quad (20)$$

if \mathcal{P} is a maximization problem (and $\min \{w_{\mathcal{I}}(x) \mid Ax \leq b\} \geq C(\mathcal{I})$ if \mathcal{P} is a minimization problem).

The *size* of the formulation is the number of inequalities in $Ax \leq b$. Finally, the (C, S) -approximate *LP formulation complexity* $\text{fc}_{\text{LP}}(\mathcal{P}, C, S)$ of \mathcal{P} is the minimal size of all its LP formulations.

The definition of SDP formulations is similar.

Definition 2.24 (SDP formulation of an optimization problem). As in Definition 2.23, let $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ be an optimization problem and C, S be real-valued functions on \mathcal{I} , the *completeness guarantee* and *soundness guarantee*. Let $\mathcal{I}^S := \{\mathcal{I} \in \mathcal{I} \mid \max \text{val}_{\mathcal{I}} \leq S(\mathcal{I})\}$ if \mathcal{P} is a maximization problem, and let $\mathcal{I}^S := \{\mathcal{I} \in \mathcal{I} \mid \min \text{val}_{\mathcal{I}} \geq S(\mathcal{I})\}$ if \mathcal{P} is a minimization problem.

A (C, S) -approximate *SDP formulation* of \mathcal{P} consists of a linear map $\mathcal{A}: \mathbb{S}^r \rightarrow \mathbb{R}^k$ and a vector $b \in \mathbb{R}^k$ (i.e., a semidefinite program $\{X \in \mathbb{S}_+^r \mid \mathcal{A}(X) = b\}$) together with the following *realizations* of \mathcal{P} :

Feasible solutions as vectors $X^s \in \mathbb{S}_+^r$ for all $s \in \mathcal{S}$ satisfying

$$\mathcal{A}(X^s) = b \quad (21)$$

i.e., the SDP $\mathcal{A}(X) = b, X \in \mathbb{S}_+^r$ is a relaxation of $\text{conv}(X^s \mid s \in \mathcal{S})$.

Instances as affine functions $w_{\mathcal{I}}: \mathbb{S}^r \rightarrow \mathbb{R}$ for all $\mathcal{I} \in \mathcal{I}^S$ satisfying

$$w_{\mathcal{I}}(X^s) = \text{val}_{\mathcal{I}}(s) \quad \text{for all } s \in \mathcal{S}, \quad (22)$$

i.e., the linearization $w_{\mathcal{I}}$ of $\text{val}_{\mathcal{I}}$ is exact on the X^s with $s \in \mathcal{S}$.

Achieving (C, S) approximation guarantee by requiring

$$\max \{w_{\mathcal{I}}(X) \mid \mathcal{A}(X^s) = b, X^s \in \mathbb{S}_+^r\} \leq C(\mathcal{I}) \quad \text{for all } \mathcal{I} \in \mathcal{I}^S, \quad (23)$$

if \mathcal{P} is a maximization problem, and the analogous inequality if \mathcal{P} is a minimization problem.

The *size* of the formulation is the dimension parameter r . Now the (C, S) -approximate *SDP formulation complexity* $\text{fc}_{\text{SDP}}(\mathcal{P}, C, S)$ of the problem \mathcal{P} is the minimal size of all its SDP formulations.

2.3.1 Slack matrix and formulation complexity

The (C, S) -approximate complexity of a maximization problem $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ is the complexity of proofs of $\text{val}_{\mathcal{I}} \leq C(\mathcal{I})$ for instances with $\max \text{val}_{\mathcal{I}} \leq S(\mathcal{I})$, and similarly for minimization problems. Formally, the proof complexity of the nonnegativity problem $\mathcal{P}_{C,S} = (\mathcal{S}, \mathcal{I}^S, C - \text{val})$ equals the (C, S) -approximate complexity of \mathcal{P} both in the LP and SDP world, as obvious from the definitions:

$$\text{fc}_{\text{LP}}(\mathcal{P}, C, S) = \text{fc}_{\text{LP}}(\mathcal{P}_{C,S}), \quad \text{fc}_{\text{SDP}}(\mathcal{P}, C, S) = \text{fc}_{\text{SDP}}(\mathcal{P}_{C,S}). \quad (24)$$

Thus the theory of nonnegativity problems from Section 2.2 immediately applies, which we formulate now explicitly for optimization problems. The material here already appeared in [BPZ15] without using nonnegativity problems and a significantly weaker reduction mechanism.

The main technical tool for establishing lower bounds on the formulation complexity of a problem is its slack matrix and its factorizations (decompositions). We start by recalling the definition of the slack matrix for optimization problems.

Definition 2.25. Let $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ be an optimization problem with *completeness guarantee* C and *soundness guarantee* S . The (C, S) -approximate slack matrix $M_{\mathcal{P}, C, S}$ is the nonnegative $\mathcal{I}^S \times \mathcal{S}$ matrix with entries

$$M_{\mathcal{P}, C, S}(\mathcal{I}, s) := \tau \cdot (C(\mathcal{I}) - \text{val}_{\mathcal{I}}(s)), \quad (25)$$

where $\tau = +1$ if \mathcal{P} is a maximization problem, and $\tau = -1$ if \mathcal{P} is a minimization problem.

Finally, we are ready to recall the factorization theorem, equating LP rank and SDP rank with LP formulation complexity and SDP formulation complexity, respectively. The notion of LP and SDP rank is recalled in Definition 2.18.

Theorem 2.26 (Factorization theorem, [BPZ15]). *Let $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ be an optimization problem with completeness guarantee C and soundness guarantee S . Then*

$$\text{f}_{\text{LP}}(\mathcal{P}, C, S) = \text{rk}_{\text{LP}} M_{\mathcal{P}, C, S}, \quad (26)$$

$$\text{f}_{\text{SDP}}(\mathcal{P}, C, S) = \text{rk}_{\text{SDP}} M_{\mathcal{P}, C, S}, \quad (27)$$

where $M_{\mathcal{P}, C, S}$ is the (C, S) -approximate slack matrix of \mathcal{P} .

Now Theorem 3.2 follows as a special case of Theorem 2.22.

2.3.2 Lasserre or SoS hierarchy

The Lasserre hierarchy, also called the Sum-of-Squares (SoS) hierarchy, is a series of SDP formulations of an optimization problem, relying on a set of base functions. The base functions are usually chosen so that the objectives $\text{val}_{\mathcal{I}}$ of instances are low-degree polynomials of the base functions. For brevity, we recall only the optimal bound obtained by the SDP formulation, using the notion of pseudoexpectation, which is essentially a feasible point of the SDP. We follow the definition of [LRST14, Page 3].

Definition 2.27 (Lasserre/SoS hierarchy).

Pseudoexpectation Let $\{f_1, \dots, f_\ell\}$ be real-valued functions with common domain \mathcal{S} . A *pseudoexpectation functional* $\tilde{\mathbb{E}}$ of level d over $\{f_1, \dots, f_\ell\}$ is a real-valued function with domain the vector space V of real-valued functions F with domain \mathcal{S} , which are polynomials in f_1, \dots, f_ℓ of degree at most d . A pseudoexpectation $\tilde{\mathbb{E}}$ is required to satisfy

Linearity For all $F_1, F_2 \in V$

$$\tilde{\mathbb{E}}(F_1 + F_2) = \tilde{\mathbb{E}}(F_1) + \tilde{\mathbb{E}}(F_2), \quad (28)$$

and for all $r \in \mathbb{R}$ and $F \in V$

$$\tilde{\mathbb{E}}(rF) = r \tilde{\mathbb{E}}(F) \quad (29)$$

Positivity $\tilde{\mathbb{E}}(F^2) \geq 0$ for all $F \in V$ with degree at most $d/2$ (so that $F^2 \in V$)

Normalization $\tilde{\mathbb{E}}(1) = 1$ for the constant function 1.

Lasserre or SoS value Given an optimization problem $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ and base functions f_1, \dots, f_ℓ defined on \mathcal{S} , the *degree d SoS value* or *round d Lasserre value* of an instance $\mathcal{I} \in \mathcal{I}$ is

$$\text{SoS}_d(\mathcal{I}) := \max_{\tilde{\mathbb{E}}: \deg \tilde{\mathbb{E}} \leq 2d} \tilde{\mathbb{E}}(\text{val}_{\mathcal{I}}). \quad (30)$$

Note that the base functions f_i might satisfy non-trivial polynomial relations, and therefore the vector space V need not be isomorphic to the vector space of *formal* low-degree polynomials in the f_i . For example, if the f_i are all 0/1-valued, which is a common case, then f_i^2 and f_i are the same elements of V . We would also like to mention that the degree or level d is not used consistently in the literature, some papers use $2d$ instead of our d . This results in a constant factor difference in the level, which is usually not significant.

For CSPs we shall use the usual set of base functions $X_{x_i=\alpha}$, the indicators that a variable x_i is assigned the value α . For graph problems, the solution set \mathcal{S} usually consists of vertex sets or edge sets. Therefore the common choice of base functions are the indicators X_v that a vertex or edge v lies in a solution. This has been used for UniqueGames in [SK10] establishing an $\omega(1)$ integrality gap for an approximate Lasserre hierarchy after a constant number of rounds.

2.4 Base hard problems

In this section we will recall the LP-hardness of the problems that will serve as the starting point in our later reductions. We start with the LP-hardness of the Matching problem with an inapproximability gap of $1 - \varepsilon/n$:

Theorem 2.28 ([BP15a], c.f., [Rot14]). *Let $n \in \mathbb{N}$ and $0 \leq \varepsilon < 1$.*

$$f_{\text{CLP}} \left(\text{Matching}(K_{2n}), \left\lfloor \frac{|V(H)|}{2} \right\rfloor + \frac{1-\varepsilon}{2}, \text{OPT}(H) \right) = 2^{\Theta(n)}, \quad (31)$$

where H is the placeholder for the instance, and the constant factor in the exponent depends on ε .

The following integrality gap was shown in [CLRS13] using the MaxCut Sherali-Adams integrality gap instances of [CMM09].

Theorem 2.29 ([CLRS13, Theorem 3.2]). *For any $\varepsilon > 0$ there are infinitely many n such that*

$$f_{\text{CLP}} \left(\text{MaxCut}(n), 1 - \varepsilon, \frac{1}{2} + \frac{\varepsilon}{6} \right) \geq n^{\Omega(\log n / \log \log n)}$$

We now recall the Lasserre integrality gap result for approximating Max- k -CSP from [BCV⁺12]. See also [BOGH⁺06, AAT05, STT07, Sch08, Tul09] for related results.

Theorem 2.30 ([BCV⁺12, Theorem 4.2]). *For $q \geq 2$, $\varepsilon, \kappa > 0$ and $\delta \geq 3/2$ and large enough n depending on $\varepsilon, \kappa, \delta$ and q , for every k, β satisfying $k \leq n^{1/2}$ and $(6q^k \ln q) / \varepsilon^2 \leq \beta \leq n^{(1-\kappa)(\delta-1)} / (10^{8(\delta-1)} k^{2\delta+0.75})$ there is a k -ary predicate $P: [q]^k \rightarrow \{0, 1\}$ and a Max- k -CSP(P) instance \mathcal{I} on alphabet $[q]$ with n variables and $m = \beta n$ constraints such that $\text{OPT}(\mathcal{I}) \leq O\left(\frac{1+\varepsilon}{q^k}\right)$, but the $\frac{m\eta}{16}$ round Lasserre relaxation for \mathcal{I} admits a perfect solution with parameter $\eta = 1 / \left(10^8 (\beta k^{2\delta+0.75})^{\frac{1}{\delta-1}}\right)$. In other words, $\text{SoS}_{\eta n/16}(\mathcal{I}) = 1$.*

The following LP-hardness for UniqueGames was shown in [LRS14] (based on [CLRS13, CMM09]):

Theorem 2.31 ([LRS14, Corollary 7.7]). *For every $q \geq 2, \delta > 0$ and $k \geq 1$ there exists a constant $c > 0$ such that for all $n \geq 1$*

$$f_{\text{CLP}} \left(\text{UniqueGames}(n, q), 1 - \delta, \frac{1}{q} + \delta \right) \geq cn^k$$

In other words there is no polynomial sized linear program that approximates UniqueGames within a factor of $1/q$.

3 Reductions with distortion

We now introduce a generalization of the affine reduction mechanism for LPs and SDPs as introduced in [BPZ15], answering an open question posed both in [BPZ15, BFPS15], leading to many new reductions that were impossible in the affine framework.

Definition 3.1 (Reduction). Let $\mathcal{P}_1 = (\mathcal{S}_1, \mathcal{I}_1, \text{val})$ and $\mathcal{P}_2 = (\mathcal{S}_2, \mathcal{I}_2, \text{val})$ be optimization problems with guarantees C_1, S_1 and C_2, S_2 , respectively. Let $\tau_1 = +1$ if \mathcal{P}_1 is a maximization problem, and $\tau_1 = -1$ if \mathcal{P}_1 is a minimization problem. Similarly, let $\tau_2 = \pm 1$ depending on whether \mathcal{P}_2 is a maximization problem or a minimization problem.

A *reduction* from \mathcal{P}_1 to \mathcal{P}_2 respecting the guarantees consists of

1. two mappings: $*$: $\mathcal{I}_1 \rightarrow \mathcal{I}_2$ and $*$: $\mathcal{S}_1 \rightarrow \mathcal{S}_2$ translating instances and feasible solutions independently;
2. two nonnegative $\mathcal{I}_1 \times \mathcal{S}_1$ matrices M_1, M_2

subject to the conditions

$$\tau_1 [C_1(\mathcal{I}_1) - \text{val}_{\mathcal{I}_1}(s_1)] = \tau_2 \left[C_2(\mathcal{I}_1^*) - \text{val}_{\mathcal{I}_1^*}(s_1^*) \right] M_1(\mathcal{I}_1, s_1) + M_2(\mathcal{I}_1, s_1) \quad (32\text{-complete})$$

$$\tau_2 \text{OPT}(\mathcal{I}_1^*) \leq \tau_2 S_2(\mathcal{I}_1^*) \quad \text{if } \tau_1 \text{OPT}(\mathcal{I}_1) \leq \tau_1 S_1(\mathcal{I}_1). \quad (32\text{-sound})$$

The matrices M_1 and M_2 provide extra freedom to add additional (valid) inequalities during the reduction. In fact, we might think of them as modeling more complex reductions. These matrices should have low computational overhead, which in our framework means LP or SDP rank, as will be obvious from the following special case of Theorem 2.22, see Section 2.3 for details.

Theorem 3.2. Let \mathcal{P}_1 and \mathcal{P}_2 be optimization problems with a reduction from \mathcal{P}_1 to \mathcal{P}_2 respecting the completeness guarantees C_1, C_2 and soundness guarantees S_1, S_2 of \mathcal{P}_1 and \mathcal{P}_2 , respectively. Then

$$\text{fc}_{\text{LP}}(\mathcal{P}_1, C_1, S_1) \leq \text{rk}_{\text{LP}} M_2 + \text{rk}_{\text{LP}} M_1 + \text{rk}_+ M_1 \cdot \text{fc}_{\text{LP}}(\mathcal{P}_2, C_2, S_2), \quad (33)$$

$$\text{fc}_{\text{SDP}}(\mathcal{P}_1, C_1, S_1) \leq \text{rk}_{\text{SDP}} M_2 + \text{rk}_{\text{SDP}} M_1 + \text{rk}_{\text{psd}} M_1 \cdot \text{fc}_{\text{SDP}}(\mathcal{P}_2, C_2, S_2), \quad (34)$$

where M_1 and M_2 are the matrices in the reduction as in Definition 3.1.

The corresponding multiplicative inapproximability factors can be obtained as usual, by taking the ratio of soundness and completeness.

4 Fractional optimization problems

A *fractional optimization problem* is an optimization problem where the objectives have the form of a fraction $\text{val}_{\mathcal{I}} = \text{val}_{\mathcal{I}}^n / \text{val}_{\mathcal{I}}^d$, such as for SparsestCut. In this case the affine space of the objective functions $\text{val}_{\mathcal{I}}$ of instances is typically not low dimensional, immediately ruling out small linear and semidefinite formulations. Nevertheless, there are examples of efficient linear programming based algorithms for such problems, however here the linear programs are used to find an optimal value of a linear combination of $\text{val}_{\mathcal{I}}^n$ and $\text{val}_{\mathcal{I}}^d$ (see e.g., [GTW13]). To be able to analyze the size of LPs or SDPs for such problems we refine the notion of formulation complexity from [BPZ15] to incorporate these types of linear programs, which reduces to the original definition with the choice of $\text{val}_{\mathcal{I}}^n = \text{val}_{\mathcal{I}}$ and $\text{val}_{\mathcal{I}}^d = 1$.

We now provide the formal definitions of linear programming and semidefinite formulations for fractional optimization problems. The idea is again that the complexity is essentially the proof complexity of $\text{val}_{\mathcal{I}} \leq$

$C(\mathcal{I})$ for instances with $\text{val}_{\mathcal{I}} \leq S(\mathcal{I})$. Formally, given a fractional optimization problem $\mathcal{P} = (\mathcal{S}, \mathfrak{I}, \text{val})$ with guarantees C, S , we study the nonnegativity problem $\mathcal{P}_{C,S} = (\mathcal{S}, \mathfrak{I}^S \times \{0, 1\}, \text{val}^*)$ with $\text{val}_{(\mathcal{I},0)}^* = C(\mathcal{I}) \text{val}_{\mathcal{I}}^d - \text{val}_{\mathcal{I}}^n$ (encoding $\text{val}_{\mathcal{I}} \leq C(\mathcal{I})$) and $\text{val}_{(\mathcal{I},1)}^* = \text{val}_{\mathcal{I}}^d$. The addition of val^d to the objective functions is for the technical reason to ensure that the objectives span the same affine space as the $\text{val}_{\mathcal{I}}^n$ and $\text{val}_{\mathcal{I}}^d$, i.e., to capture the affineness of these functions. This is not expected to significantly affect the complexity of the resulting problem, as the $\text{val}_{\mathcal{I}}^d$ in interesting applications are usually a positive linear combination of a small number of nonnegative functions.

As a special case of Section 2.2 we obtain the following setup for fractional optimization problems. Note that when \mathcal{P} is a fractional optimization problem with $\text{val}^d = 1$, then \mathcal{P} is an optimization problem and Definitions 4.1 and 4.2 are equivalent to Definitions 2.23 and 2.24, as we will see now.

Definition 4.1 (LP formulation of a fractional optimization problem). Let $\mathcal{P} = (\mathcal{S}, \mathfrak{I}, \text{val})$ be a fractional optimization problem and let C, S be two real valued functions on \mathfrak{I} called *completeness guarantee* and *soundness guarantee* respectively. Let $\mathfrak{I}^S := \{\mathcal{I} \in \mathfrak{I} \mid \max \text{val}_{\mathcal{I}} \leq S(\mathcal{I})\}$ when \mathcal{P} is a maximization problem and $\mathfrak{I}^S := \{\mathcal{I} \in \mathfrak{I} \mid \min \text{val}_{\mathcal{I}} \geq S(\mathcal{I})\}$ if \mathcal{P} is a minimization problem.

A (C, S) -approximate LP formulation for the problem \mathcal{P} consists of a linear program $Ax \leq b$ with $x \in \mathbb{R}^r$ for some r and the following realizations:

Feasible solutions as vectors $x^s \in \mathbb{R}^r$ for every $s \in \mathcal{S}$ satisfying

$$Ax^s \leq b \quad \text{for all } s \in \mathcal{S},$$

i.e., $Ax \leq b$ is a relaxation of $\text{conv}(x^s \mid s \in \mathcal{S})$.

Instances as a pair of affine functions $w_{\mathcal{I}}^n, w_{\mathcal{I}}^d: \mathbb{R}^r \rightarrow \mathbb{R}$ for all $\mathcal{I} \in \mathfrak{I}^S$ satisfying

$$\begin{aligned} w_{\mathcal{I}}^n(x^s) &= \text{val}_{\mathcal{I}}^n(s) \\ w_{\mathcal{I}}^d(x^s) &= \text{val}_{\mathcal{I}}^d(s) \end{aligned}$$

for every $s \in \mathcal{S}$. In other words the linearizations $w_{\mathcal{I}}^n, w_{\mathcal{I}}^d$ are required to be *exact* on all x^s for $s \in \mathcal{S}$.

Achieving (C, S) approximation guarantee requiring the following for every $\mathcal{I} \in \mathfrak{I}^S$

$$Ax \leq b \Rightarrow \begin{cases} w_{\mathcal{I}}^d(x) \geq 0 \\ w_{\mathcal{I}}^n(x) \leq C(\mathcal{I})w_{\mathcal{I}}^d(x) \end{cases}$$

if \mathcal{P} is a maximization problem and

$$Ax \leq b \Rightarrow \begin{cases} w_{\mathcal{I}}^d(x) \geq 0 \\ w_{\mathcal{I}}^n(x) \geq C(\mathcal{I})w_{\mathcal{I}}^d(x) \end{cases}$$

if \mathcal{P} is a minimization problem. In other words we can derive the nonnegativity of $w_{\mathcal{I}}^d$ and the approximation guarantee $C(\mathcal{I})$ from the set of inequalities in $Ax \leq b$.

The *size* of the formulation is the number of inequalities in $Ax \leq b$. Finally, the (C, S) -approximate LP formulation complexity $\text{fc}_{\text{LP}}(\mathcal{P}, C, S)$ of \mathcal{P} is the minimal size of all its LP formulations.

SDP formulations for fractional optimization problems are defined similarly.

Definition 4.2 (SDP formulation of fractional optimization problem). Let $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ be a fractional optimization problem and let $C, S: \mathcal{I} \rightarrow \mathbb{R}_{\geq 0}$ be the *completeness guarantee* and the *soundness guarantee* respectively. Let $\mathcal{I}^S := \{\mathcal{I} \in \mathcal{I} \mid \max \text{val}_{\mathcal{I}} \leq S(\mathcal{I})\}$ when \mathcal{P} is a maximization problem and $\mathcal{I}^S := \{\mathcal{I} \in \mathcal{I} \mid \min \text{val}_{\mathcal{I}} \geq S(\mathcal{I})\}$ if \mathcal{P} is a minimization problem.

A (C, S) -approximate SDP formulation of \mathcal{P} consists of a linear map $\mathcal{A}: \mathbb{S}^r \rightarrow \mathbb{R}^k$ together with a vector $b \in \mathbb{R}^k$ (i.e., a semidefinite program $\{X \in \mathbb{S}_+^r \mid \mathcal{A}(X) = b\}$) and the following realizations of \mathcal{P} :

Feasible solutions as vectors $X^s \in \mathbb{S}_+^r$ for every $s \in \mathcal{S}$ satisfying

$$\mathcal{A}(X^s) = b \quad \text{for every } s \in \mathcal{S},$$

i.e. $\mathcal{A}(X) = b, X \in \mathbb{S}_+^r$ is a relaxation of $\text{conv}(X^s \mid s \in \mathcal{S})$.

Instances as a pair of affine functions $w_{\mathcal{I}}^n, w_{\mathcal{I}}^d: \mathbb{S}^r \rightarrow \mathbb{R}_{\geq 0}$ for every $\mathcal{I} \in \mathcal{I}^S$ satisfying

$$w_{\mathcal{I}}^n(X^s) = \text{val}_{\mathcal{I}}^n(s)$$

$$w_{\mathcal{I}}^d(X^s) = \text{val}_{\mathcal{I}}^d(s)$$

for every $s \in \mathcal{S}$. In other words the linearizations $w_{\mathcal{I}}^n, w_{\mathcal{I}}^d$ are required to be *exact* on all X^s for $s \in \mathcal{S}$.

Achieving (C, S) approximation guarantee requiring the following for every $\mathcal{I} \in \mathcal{I}^S$

$$\mathcal{A}(X) = b \Rightarrow \begin{cases} w_{\mathcal{I}}^d(X) \geq 0 \\ w_{\mathcal{I}}^n(X) \leq C(\mathcal{I})w_{\mathcal{I}}^d(X) \end{cases}$$

if \mathcal{P} is a maximization problem and

$$\mathcal{A}(X) = b \Rightarrow \begin{cases} w_{\mathcal{I}}^d(X) \geq 0 \\ w_{\mathcal{I}}^n(X) \leq C(\mathcal{I})w_{\mathcal{I}}^d(X) \end{cases}$$

if \mathcal{P} is a minimization problem.

The *size* of the formulation is given by the dimension r . The (C, S) -approximate SDP formulation complexity $\text{fc}_{\text{SDP}}(\mathcal{P}, C, S)$ of the problem \mathcal{P} is the minimal size of all its SDP formulations.

The slack matrix for fractional problems plays the same role as for non-fractional problems, with the twist that we factorize the denominator and numerator separately. This allows us to overcome the high dimensionality of the space spanned by the actual ratios.

Definition 4.3. Let $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ be a fractional optimization problem with *completeness guarantee* C and *soundness guarantee* S . The (C, S) -approximate slack matrix $M_{\mathcal{P}, C, S}$ is the nonnegative $2\mathcal{I}^S \times \mathcal{S}$ matrix of the form

$$M_{\mathcal{P}, C, S} = \begin{bmatrix} M_{\mathcal{P}, C, S}^{(d)} \\ M_{\mathcal{P}, C, S}^{(n)} \end{bmatrix}$$

where $M_{\mathcal{P}, C, S}^{(d)}, M_{\mathcal{P}, C, S}^{(n)}$ are nonnegative $\mathcal{I}^S \times \mathcal{S}$ matrices with entries

$$M_{\mathcal{P}, C, S}^{(d)}(\mathcal{I}, s) := \text{val}_{\mathcal{I}}^d(s)$$

$$M_{\mathcal{P}, C, S}^{(n)}(\mathcal{I}, s) := \tau \left(C(\mathcal{I}) \text{val}_{\mathcal{I}}^d(s) - \text{val}_{\mathcal{I}}^n(s) \right)$$

where $\tau = +1$ if \mathcal{P} is a maximization problem and $\tau = -1$ if \mathcal{P} is a minimization problem.

We are now ready to obtain the factorization theorem for the class of fractional optimization problems, as a special case of Theorem 2.20:

Theorem 4.4 (Factorization theorem for fractional optimization problems). *Let $\mathcal{P} = (\mathcal{S}, \mathcal{I}, \text{val})$ be a fractional optimization problem with completeness guarantee C and soundness guarantee S . Then*

$$\begin{aligned} \text{f}_{\text{CLP}}(\mathcal{P}, C, S) &= \text{rk}_{\text{LP}} M_{(\mathcal{P}, C, S)}, \\ \text{f}_{\text{CDP}}(\mathcal{P}, C, S) &= \text{rk}_{\text{SDP}} M_{(\mathcal{P}, C, S)} \end{aligned}$$

where $M_{(\mathcal{P}, C, S)}$ is the (C, S) -approximate slack matrix of \mathcal{P} .

Now Theorem 4.6 arises as a special case of Theorem 2.22.

4.1 Reduction between fractional problems

Reductions for fractional optimization problems are completely analogous to the non-fractional case:

Definition 4.5 (Reduction). Let $\mathcal{P}_1 = (\mathcal{S}_1, \mathcal{I}_1, \text{val})$ and $\mathcal{P}_2 = (\mathcal{S}_2, \mathcal{I}_2, \text{val})$ be fractional optimization problems with guarantees C_1, S_1 and C_2, S_2 , respectively. Let $\tau_1 = +1$ if \mathcal{P}_1 is a maximization problem, and $\tau_1 = -1$ if \mathcal{P}_1 is a minimization problem. Similarly, let $\tau_2 = \pm 1$ depending on whether \mathcal{P}_2 is a maximization problem or a minimization problem.

A reduction from \mathcal{P}_1 to \mathcal{P}_2 respecting the guarantees consists of

1. two mappings: $*$: $\mathcal{I}_1 \rightarrow \mathcal{I}_2$ and $*$: $\mathcal{S}_1 \rightarrow \mathcal{S}_2$ translating instances and feasible solutions independently;
2. four nonnegative $\mathcal{I}_1 \times \mathcal{S}_1$ matrices $M_1^{(n)}, M_1^{(d)}, M_2^{(n)}, M_2^{(d)}$

subject to the conditions

$$\tau_1 \left[C_1(\mathcal{I}_1) \text{val}_{\mathcal{I}_1}^d(s_1) - \text{val}_{\mathcal{I}_1}^n(s_1) \right] = \tau_2 \left[C_2(\mathcal{I}_1^*) \text{val}_{\mathcal{I}_1^*}^d(s_1^*) - \text{val}_{\mathcal{I}_1^*}^n(s_1^*) \right] M_1^{(n)}(\mathcal{I}_1, s_1) + M_2^{(n)}(\mathcal{I}_1, s_1) \quad (35\text{-complete})$$

$$\text{val}_{\mathcal{I}_1}^d(s_1) = \text{val}_{\mathcal{I}_1^*}^d(s_1^*) \cdot M_1^{(d)}(\mathcal{I}_1, s_1) + M_2^{(d)}(\mathcal{I}_1, s_1) \quad (35\text{-denominator})$$

$$\tau_2 \text{OPT}(\mathcal{I}_1^*) \leq \tau_2 S_2(\mathcal{I}_1^*) \quad \text{if } \tau_1 \text{OPT}(\mathcal{I}_1) \leq \tau_1 S_1(\mathcal{I}_1). \quad (35\text{-sound})$$

As the val^d are supposed to have a small proof, the matrices $M_1^{(d)}$ and $M_2^{(d)}$ are not supposed to significantly influence the strength of the reduction even with the trivial choice $M_1^{(d)} = 0$ and $M_2^{(d)}(\mathcal{I}_1, s_1) = \text{val}_{\mathcal{I}_1}^d(s_1)$. However, as in the non-fractional case, the complexity of $M_1^{(n)}$ and $M_2^{(n)}$ could have a major influence on the strength of the reduction. The reduction theorem is a special case of Theorem 2.22, see Section 4:

Theorem 4.6. *Let \mathcal{P}_1 and \mathcal{P}_2 be optimization problems with a reduction from \mathcal{P}_1 to \mathcal{P}_2 . Then*

$$\text{f}_{\text{CLP}}(\mathcal{P}_1, C_1, S_1) \leq \text{rk}_{\text{LP}} \begin{bmatrix} M_2^{(n)} \\ M_2^{(d)} \end{bmatrix} + \text{rk}_{\text{LP}} \begin{bmatrix} M_1^{(n)} \\ M_1^{(d)} \end{bmatrix} + \text{rk}_+ \begin{bmatrix} M_1^{(n)} \\ M_1^{(d)} \end{bmatrix} \cdot \text{f}_{\text{CLP}}(\mathcal{P}_2, C_2, S_2), \quad (36)$$

$$\text{f}_{\text{CDP}}(\mathcal{P}_1, C_1, S_1) \leq \text{rk}_{\text{SDP}} \begin{bmatrix} M_2^{(n)} \\ M_2^{(d)} \end{bmatrix} + \text{rk}_{\text{SDP}} \begin{bmatrix} M_1^{(n)} \\ M_1^{(d)} \end{bmatrix} + \text{rk}_{\text{psd}} \begin{bmatrix} M_1^{(n)} \\ M_1^{(d)} \end{bmatrix} \cdot \text{f}_{\text{CDP}}(\mathcal{P}_2, C_2, S_2), \quad (37)$$

where $M_1^{(n)}, M_1^{(d)}, M_2^{(n)}$, and $M_2^{(d)}$ are the matrices in the reduction as in Definition 4.5.

5 A simple example: Matching over 3-regular graphs has no small LPs

We now show that the Matching problem even over 3-regular graphs does not admit a small LP formulation. This has been an open question of various researchers, given that the Matching problem admits polynomial-size LPs for many classes of sparse graphs, like bounded treewidth, planar (and bounded genus) graphs [Bar93, Ger91, KKT15]. We also show that for graphs of bounded degree 3, the Matching problem does not admit fully-polynomial size relaxation schemes, the linear programming equivalent of FPTAS, see [BP15b, BP15a] for details on these schemes.

Theorem 5.1. *Let $n \in \mathbb{N}$ and $0 \leq \varepsilon < 1$. There exists a 3-regular graph D_{2n} with $2n(2n - 1)$ vertices, so that*

$$\text{f}_{\text{LP}} \left(\text{Matching}(D_{2n}), \left\lfloor \frac{|V(H)|}{2} \right\rfloor + \frac{1 - \varepsilon}{2}, \text{OPT}(H) \right) = 2^{\Omega(\sqrt{|V(D_{2n})|})}, \quad (38)$$

where H is the placeholder for an instance, and the constant factor in the exponent depends on ε . In particular, $\text{Matching}(D_{2n})$ is LP-hard with an inapproximability factor of $1 - \varepsilon / |V(D_{2n})|$.

Proof. As usual, the inapproximability factor simply arises as the smallest factor $\text{OPT}(H) / (\lfloor |V(H)| / 2 \rfloor + (1 - \varepsilon) / 2)$ of the soundness and completeness guarantees.

The proof is a simple application of the reduction framework. In fact, it suffices to use the affine framework of [BPZ15]. We will reduce from the perfect matching problem $\text{Matching}(K_{2n})$ as given in Definition 2.10.

We first construct our target graph D_{2n} as follows, see Figure 1:

1. For every vertex v of K_{2n} we consider a cycle C^v of length $2n - 1$. We denote the vertices of C^v by $[v, u]$, where $v, u \in V$ and $v \neq u$.
2. The graph D_{2n} is the disjoint union of the C^v for $v \in V$ together with the following additional edges: an edge $([v, u], [u, v])$ for every $(u, v) \in E$.

Thus D_{2n} has a total of $2n(2n - 1)$ vertices. This completes the definition of the graph D_{2n} , which is




Figure 1: The graph D_{2n} for $n = 2$ in the reduction to 3-regular Matching.

obviously 3-regular. (There is some ambiguity regarding the order of vertices in the cycles C^v , but this does not affect the argument below.) Now we define the reduction from $\text{Matching}(K_{2n})$ to $\text{Matching}(D_{2n})$.

We first map the instances. Let H be a subgraph of K_{2n} . Its image H^* under the reduction is the union of the C^v for $v \in H$ together with the edges $([u, v], [v, u])$ for $\{u, v\} \in E(H)$.

Now let M be a perfect matching in K_{2n} . We define M^* by naturally extending it to a perfect matching in D_{2n} . For every edge $e = \{u, v\} \in M$ in the matching, the edges $([u, v], [v, u]) \in D_{2n}$ form a matching containing exactly one vertex from every cycle C^v . We choose M to be the unique extension of this matching to a perfect matching by adding edges from the cycles C^v .

We obviously have the following relationship between the objective values:

$$\begin{aligned}\text{val}_{H^*}^{D_{2n}}(M^*) &= |M^* \cap E(H^*)| = |V(H)| \cdot (n-1) + |M \cap V(H)| \\ &= |V(H)| \cdot (n-1) + \text{val}_H^{K_{2n}}(M),\end{aligned}\tag{39}$$

providing immediately the completeness of the reduction:

$$\begin{aligned}\left\lfloor \frac{|V(H)|}{2} \right\rfloor + \frac{1-\varepsilon}{2} - \text{val}_H^{K_{2n}}(M) &= \left\lfloor \frac{|V(H)| \cdot (2n-1)}{2} \right\rfloor + \frac{1-\varepsilon}{2} - \text{val}_{H^*}^{D_{2n}}(M^*) \\ &= \left\lfloor \frac{|V(H^*)|}{2} \right\rfloor + \frac{1-\varepsilon}{2} - \text{val}_{H^*}^{D_{2n}}(M^*).\end{aligned}\tag{40}$$

The soundness of the reduction is immediate, as the soundness guarantee is the optimal value. \square

It is an interesting open problem, whether there exists a family of bounded-degree graphs G_n on n vertices so that the lower bound in Theorem 5.1 can be strengthened to $2^{\Omega(n)}$.

6 BalancedSeparator and SparsestCut

The SparsestCut problem is a high-profile problem that received considerable attention in the past. It is known that SparsestCut with general demands can be approximated within a factor of $O(\sqrt{\log n \log \log n})$ [ALN08] and that the standard SDP has an integrality gap of $(\log n)^{\Omega(1)}$ [CKN09]. The BalancedSeparator problem is a related problem which often arises in connection to the SparsestCut problem (see Definition 2.4). The main result of this section will be to show that the SparsestCut and BalancedSeparator problems cannot be approximated well by small LPs and SDPs by using the new reduction mechanism from Section 4.1. In the case of the SparsestCut problem our result holds even if the supply graph has bounded treewidth, with the lower bound matching the upper bound in [GTW13] in the LP case. The results are *unconditional* LP/SDP analogues to [CKK⁺06], however for a different regime. In the case of the BalancedSeparator problem our result holds even if the demand graph has bounded treewidth.

The SparsestCut problem is a fractional optimization problem: we extend Definition 2.3 via

$$\text{val}_{\mathcal{I}}^c(s) := \sum_{i \in s, j \notin s} c(i, j), \quad \text{val}_{\mathcal{I}}^d(s) := \sum_{i \in s, j \notin s} d(i, j)\tag{41}$$

for any vertex set s and any instance \mathcal{I} with capacity c and demand d .

Theorem 6.1 (LP/SDP hardness for SparsestCut, $\text{tw}(\text{supply}) = O(1)$). *For any $\varepsilon \in (0, 1)$ there are $\eta_{LP} > 0$ and $\eta_{SDP} > 0$ such that for every large enough n the following hold*

$$\begin{aligned}\text{fc}_{LP}(\text{SparsestCut}(n, 2), \eta_{LP}(1 + \varepsilon), \eta_{LP}(2 - \varepsilon)) &\geq n^{\Omega(\log n / \log \log n)}, \\ \text{fc}_{SDP}\left(\text{SparsestCut}(n, 2), \eta_{SDP}\left(1 + \frac{4\varepsilon}{5}\right), \eta_{SDP}\left(\frac{16}{15} - \varepsilon\right)\right) &\geq n^{\Omega(\log n / \log \log n)}.\end{aligned}$$

In other words SparsestCut($n, 2$) is LP-hard with an inapproximability factor of $2 - \varepsilon$, and SDP-hard with an inapproximability factor of $\frac{16}{15} - O(\varepsilon)$.

A complementary reduction proves the hardness of approximating BalancedSeparator where the demand graph has constant treewidth. Note that we only have an inapproximability result for LPs in this case since the reduction is from UniqueGames for which we do not yet know of any SDP hardness result.

Theorem 6.2 (LP-hardness for BalancedSeparator). *For any constant $c_1 \geq 1$ there is another constant $c_2 \geq 1$ such that for all n there is a demand function $d: E(K_n) \rightarrow \mathbb{R}_{\geq 0}$ satisfying $\text{tw}([n]_d) \leq c_2$ so that BalancedSeparator(n, d) is LP-hard with an inapproximability factor of c_1 .*

6.1 SparsestCut with bounded treewidth supply graph

In this section we show that the SparsestCut problem over supply graphs with treewidth 2 cannot be approximated up to a factor of 2 by any polynomial sized LP and up to a factor of $\frac{16}{15}$ by any polynomial sized SDP, i.e., Theorem 6.1.

We use the reduction from [GTW13], reducing MaxCut to SparsestCut. Given an instance \mathcal{I} of MaxCut(n) we first construct the instance \mathcal{I}^* on vertex set $V = \{u, v\} \cup [n]$ where u and v are two special vertices. Let us denote the degree of a vertex i in \mathcal{I} by $\deg(i)$ and let $m := \frac{1}{2} \sum_{i=1}^n \deg(i)$ be the total number of edges in \mathcal{I} . We define the capacity function $c: V \times V \rightarrow \mathbb{R}_{\geq 0}$ as

$$c(i, j) := \begin{cases} \frac{\deg(i)}{m} & \text{if } j = u, i \neq v \text{ or } j = v, i \neq u \\ 0 & \text{otherwise.} \end{cases}$$

Note that the supply graph has treewidth at most 2 being a copy of $K_{2,n}$. The demand function $d: V \times V \rightarrow \mathbb{R}_{\geq 0}$ is defined as

$$d(i, j) := \begin{cases} \frac{2}{m} & \text{if } \{i, j\} \in E(\mathcal{I}) \\ 0 & \text{otherwise.} \end{cases}$$

We map a solution s to MaxCut(n) to the cut $s^* := s \cup \{u\}$ of SparsestCut($n+2, 2$).

We remind the reader of the powering operation from [GTW13] to handle the case of unbalanced and non u - v cuts. It successively adds for every edge of \mathcal{I}^* a copy of itself, scaling both the capacities and demands by the capacity of the edge. After l rounds, we obtain an instance \mathcal{I}_l^* on a fixed set of $O(N^{2l})$ vertices, and similarly the cuts s^* extend naturally to cuts s_l^* on these vertices, independent of the instance \mathcal{I} . We provide a formal definition of the powering operation below, for any general instance \mathcal{I}_1 and general solution s_1 of SparsestCut.

Definition 6.3 (Powering instances). The instances of SparsestCut(N_1) are $G_1 := K_{N_1}$ with capacity function c_1 and demand function d_1 . Let u and v be two distinguished vertices of G_1 . We construct a sequence $\{G_l\}_l$ of graphs with distinguished vertices u and v recursively as follows. The graph G_l is obtained by replacing every edge $\{x, y\}$ of G_1 by a copy of G_{l-1} . Let us denote by $(\{x, y\}, w)$ the copy of vertex w of G_{l-1} . We identify the vertices $(\{x, y\}, u)$ and $(\{x, y\}, v)$ with x and y . There are two ways to do so for every edge and we can pick either, arbitrarily. Obviously, G_l has $N_l := \sum_{i=1}^{l-1} \binom{N}{2}^i (N-2) + 2$ many vertices. Given a base instance \mathcal{I}_1 of SparsestCut(N_1) we will construct a sequence of instances $\{\mathcal{I}_l\}_l$ of SparsestCut(N_l) recursively as follows. Let the capacity and demand function of \mathcal{I}_{l-1} be c_{l-1} and d_{l-1} respectively. The capacity of edges not in G_l will be 0. Any edge e of G_l has the unique form $\{(\{x, y\}, p), (\{x, y\}, q)\}$ for an edge $\{x, y\}$ of G_1 and an edge $\{p, q\}$ of G_{l-1} . We define $c_l(e) := c_{l-1}(p, q) \cdot c_1(x, y)$. If e is not the edge $\{x, y\}$ then let $d_l(e) := d_{l-1}(p, q) \cdot c_1(x, y)$. The edge $\{x, y\}$ takes the demand from G_1 in addition, therefore we define $d_l(x, y) := d_{l-1}(u, v) \cdot c_1(x, y) + d_1(x, y)$.

We recall here the following easy observation that relates the treewidth of the supply graph of \mathcal{I}_1 to the treewidth of the supply graph of \mathcal{I}_l .

Lemma 6.4 ([GTW13, Observation 4.4]). *If the treewidth of the supply graph of \mathcal{I}_1 is at most k , then the treewidth of the supply graph of \mathcal{I}_l is also at most k .*

Corresponding to powering instances, we can also recursively construct solutions to SparsestCut(N_l) starting from a solution s_1 of SparsestCut(N_1).

Definition 6.5 (Powering solutions separating u and v). Given a base solution s_1 of $\text{SparsestCut}(N_1)$ we construct a solution s_l for $\text{SparsestCut}(N_l)$ recursively as follows. The solution s_l coincides with s_1 on the vertices of G_1 . On the copy of G_{l-1} for an edge $\{x, y\}$ of G_1 we define s_l as follows. If $s_1(x) = s_1(y)$ then let $s_l(\{x, y\}, z) := s_1(x) = s_1(y)$ for all vertex z of G_{l-1} , so as s_l cuts no edges in the copy of G_{l-1} . If $s_1(x) \neq s_1(y)$ then we define s_l so that the edges it cuts in the $\{x, y\}$ -copy of G_{l-1} are exactly the copies of edges cut by s_{l-1} in G_{l-1} . More precisely, let $(\{x, y\}, u)$ be identified with x and $(\{x, y\}, v)$ with y . If $s_1(x) = s_{l-1}(u)$ then we let $s_l(\{x, y\}, z) := s_{l-1}(z)$, otherwise we let $s_l(\{x, y\}, z) := -s_{l-1}(z)$.

We now define the actual reduction. We construct a sequence of instances $\{\mathcal{I}_1^*, \mathcal{I}_2^*, \dots, \mathcal{I}_l^*\}$ where \mathcal{I}_l^* is obtained as in Definition 6.3 by applying the powering operation to the base instance $\mathcal{I}_1^* = \mathcal{I}^*$ and where $N_1 := n + 2, c_1 := c$ and $d_1 := d$. Note that by Lemma 6.4, the treewidth of the supply graph of \mathcal{I}_l^* is at most 2. We also construct a sequence of solutions $\{s_1^*, \dots, s_l^*\}$ where s_l^* is obtained as in Definition 6.5 by applying the powering operation to the base solution $s_1^* = s^*$. The final reduction maps the instance \mathcal{I} and solution s of $\text{MaxCut}(n)$ to the instance \mathcal{I}_l^* and solution s_l^* of $\text{SparsestCut}(N_l, 2)$ respectively. Completeness and soundness follows from [GTW13].

Lemma 6.6 (Completeness, [GTW13, Claim 4.2]). *Let \mathcal{I} be an instance and s be a solution of $\text{MaxCut}(n)$, and let their image be the instance \mathcal{I}_l^* and solution s_l^* of $\text{SparsestCut}(N_l, 2)$, respectively. Then the following holds*

$$\text{val}_{\mathcal{I}_l^*}^n(s_l^*) = 1, \quad \text{val}_{\mathcal{I}_l^*}^d(s_l^*) = l \text{val}_{\mathcal{I}}(s).$$

Lemma 6.7 (Soundness, [GTW13, Lemmas 4.3 and 4.7]). *Let \mathcal{I} be an instance of $\text{MaxCut}(n)$ and let \mathcal{I}_l^* be the instance of $\text{SparsestCut}(N_l, 2)$ it is mapped to. Then the instance \mathcal{I}_l^* has the following lower bound on its optimum (the number of edges of \mathcal{I} scales the MaxCut value between 0 and 1).*

$$\text{OPT}(\mathcal{I}_l^*) \geq \frac{1}{1 + (l-1) \text{OPT}(\mathcal{I}) / |E(\mathcal{I})|}.$$

Using the reduction framework of Section 4.1 we now prove the main theorem of this section about the LP and SDP inapproximability of SparsestCut .

Proof of Theorem 6.1. This is a simple application of Lemmas 6.6 and 6.7 using Theorem 4.6 with matrices $M_1^{(n)}(\mathcal{I}_1, s_1) := C_1(\mathcal{I}_1), M_2^{(n)}(\mathcal{I}_1, s_1) := 0, M_1^{(d)}(\mathcal{I}_1, s_1) := 0, M_2^{(d)}(\mathcal{I}_1, s_1) := 1$. Hardness of the base problem MaxCut is provided by Theorems 2.29 and 7.1, and leads to $\eta_{\text{LP}} = \frac{5\epsilon}{3-\epsilon}$ and $\eta_{\text{SDP}} = \frac{3\epsilon}{1-4\epsilon}$. \square

6.2 BalancedSeparator with bounded-treewidth demand graph

In this section we show that the BalancedSeparator problem cannot be approximated within any constant factor with small LPs even when the demand graph has constant treewidth:

Theorem 6.8 (LP-hardness for BalancedSeparator). *(Theorem 6.2 restated) For any constant $c_1 \geq 1$ there is another constant $c_2 \geq 1$ such that for all n there is a demand function $d: E(K_n) \rightarrow \mathbb{R}_{\geq 0}$ satisfying $\text{tw}([n]_d) \leq c_2$ so that $\text{BalancedSeparator}(n, d)$ is LP-hard with an inapproximability factor of c_1 .*

We will reduce the $\text{UniqueGames}(n, q)$ problem to the $\text{BalancedSeparator}(2^q n, d)$ problem for a fixed demand function d to be defined below. We reuse the reduction from [KV15, Section 11.1]. A bijection $\pi: [q] \rightarrow [q]$ acts on strings $\{-1, 1\}^q$ in the natural way, i.e., $\pi(x)_i := x_{\pi(i)}$. For any parameter $p \in [0, 1]$, we denote by $\mathbf{x} \in_p \{-1, 1\}^q$ a random string where each coordinate x_i of \mathbf{x} is -1 with probability p and 1 with probability $1 - p$. For a string $x \in \{-1, 1\}^q$ we define $x_+ := |\{i \mid x_i = 1\}|$ and $x_- := |\{i \mid x_i = -1\}|$. For a pair of strings $x, y \in \{-1, 1\}^q$ we denote by xy the string in $\{-1, 1\}^q$ formed by the

coordinate-wise product of x and y , i.e., $(xy)_i := x_i y_i$ for $i \in [n]$. We are now ready to proceed with the reduction.

Given an instance $\mathcal{I} = \mathcal{I}(w, \pi)$ of $\text{UniqueGames}(n, q)$ we construct the instance \mathcal{I}^* of $\text{BalancedSeparator}(2^q n, d)$. Let ε be a parameter to be chosen later. The vertex set V of \mathcal{I}^* is defined as $V := \{(x, i) \mid i \in [n], x \in \{-1, 1\}^q\}$ so that $|V| = 2^q n$. Let $W := \sum_{\{i, j\} \in E(K_n)} w(i, j)$ denote the total weight of the $\text{UniqueGames}(n, q)$ instance \mathcal{I} . For every $i, j \in [n]$ and $x, y \in \{-1, 1\}^q$ there is an undirected edge $\{(x, i), (y, j)\}$ in \mathcal{I}^* of capacity $c((x, i), (y, j))$ which is defined as

$$c((x, i), (y, j)) := \frac{w(i, j) \varepsilon^{\left(\pi_{i,j}(x)y\right)_-} (1 - \varepsilon)^{\left(\pi_{i,j}(x)y\right)_+}}{2^q W}.$$

The demand function $d((x, i), (y, j))$ is defined for an unordered pair of vertices $\{(x, i), (y, j)\}$ as

$$d((x, i), (y, j)) := \begin{cases} \frac{1}{2^{2q-1}n} & \text{if } i = j \\ 0 & \text{otherwise} \end{cases}$$

so that the total demand D is 1. Note that the demand graph $[2^q n]_d$ is a disjoint union of cliques of size 2^q and so $\text{tw}([2^q n]_d) = 2^q - 1 = O(1)$. Given a solution s of $\text{UniqueGames}(n, q)$ we map it to the solution s^* of $\text{BalancedSeparator}(2^q n, d)$ defined as $s^* := \{(x, i) \mid x_{s(i)} = 1\}$. Note that the total demand cut by s^* is $\frac{1}{2} = \frac{D}{2} > \frac{D}{4}$ since for every solution $s(i) \in [q]$ there are exactly 2^{q-1} strings in $\{-1, 1\}^q$ that have their $s(i)^{\text{th}}$ bit set to 1 and 2^{q-1} strings have their $s(i)^{\text{th}}$ bit set to -1 . Thus s^* is a valid solution to the $\text{BalancedSeparator}(2^q n, d)$ problem and moreover is independent of the instance \mathcal{I}^* . We are now ready to show that this reduction satisfies completeness.

Lemma 6.9 (Completeness). *Let \mathcal{I} and s be an instance and a solution respectively of $\text{UniqueGames}(n, q)$. Let \mathcal{I}^* and s^* be the instance and solution of $\text{BalancedSeparator}(2^q n, d)$ obtained from the reduction. Then*

$$\frac{1}{2} - \text{val}_{\mathcal{I}^*}(s^*) = \left(\frac{1}{2} - \varepsilon\right) \text{val}_{\mathcal{I}}(s)$$

Proof. Let us sample a random edge (\mathbf{i}, \mathbf{j}) from the $\text{UniqueGames}(n, q)$ instance \mathcal{I} with probabilities proportional to $w(i, j)$ (i.e., $\mathbb{P}[\mathbf{i} = i, \mathbf{j} = j] = w(i, j)/W$), and independently sample $\mathbf{x} \in_{1/2} \{-1, 1\}^q$ and $\mathbf{z} \in_{\varepsilon} \{-1, 1\}^q$. Let $\mathbf{y} := \pi_{\mathbf{i}, \mathbf{j}}(\mathbf{x})\mathbf{z}$.

The claim follows by computing the probability of $\mathbf{x}_{s(\mathbf{i})} = \mathbf{y}_{s(\mathbf{j})}$ in two different ways.

On the one hand, for a fixed edge (\mathbf{i}, \mathbf{j}) of \mathcal{I} , depending on whether the edge is correctly labelled, we have

$$\mathbb{P}[\mathbf{x}_{s(\mathbf{i})} = \mathbf{y}_{s(\mathbf{j})} \mid \mathbf{i} = i, \mathbf{j} = j, s(i) = \pi_{i,j}(s(j))] = \mathbb{P}[\mathbf{z}_{s(j)} = 1 \mid \mathbf{i} = i, \mathbf{j} = j, s(i) = \pi_{i,j}(s(j))] = 1 - \varepsilon, \quad (42)$$

$$\mathbb{P}[\mathbf{x}_{s(\mathbf{i})} = \mathbf{y}_{s(\mathbf{j})} \mid \mathbf{i} = i, \mathbf{j} = j, s(i) \neq \pi_{i,j}(s(j))] = \mathbb{P}[\mathbf{x}_{s(i)} = \mathbf{y}_{s(j)} \mid \mathbf{i} = i, \mathbf{j} = j, s(i) \neq \pi_{i,j}(s(j))] = \frac{1}{2}. \quad (43)$$

Note that in the latter case $\mathbf{x}_{s(i)}$ and $\mathbf{y}_{s(j)}$ are independent uniform binary variables. Hence

$$\mathbb{P}[s(\mathbf{i}) = \pi_{\mathbf{i}, \mathbf{j}}(s(\mathbf{j})), \mathbf{x}_{s(\mathbf{i})} = \mathbf{y}_{s(\mathbf{j})}] = (1 - \varepsilon) \text{val}_{\mathcal{I}}(s), \quad (44)$$

$$\mathbb{P}[s(\mathbf{i}) \neq \pi_{\mathbf{i}, \mathbf{j}}(s(\mathbf{j})), \mathbf{x}_{s(\mathbf{i})} = \mathbf{y}_{s(\mathbf{j})}] = \mathbb{P}[\mathbf{x}_{s(i)} = \mathbf{y}_{s(j)} \mid \mathbf{i} = i, \mathbf{j} = j, s(i) \neq \pi_{i,j}(s(j))] \quad (45)$$

$$= \frac{1 - \text{val}_{\mathcal{I}}(s)}{2}, \quad (46)$$

leading to

$$\mathbb{P} \left[\mathbf{x}_{s(i)} = \mathbf{y}_{s(j)} \right] = \frac{1}{2} + \left(\frac{1}{2} - \varepsilon \right) \text{val}_{\mathcal{I}}(s). \quad (47)$$

On the other hand, note that $((\mathbf{x}, \mathbf{i}), (\mathbf{y}, \mathbf{j}))$ is a random edge from \mathcal{I}^* with distribution given by the weights $c((x, i), (y, j))$, (i.e., $\mathbb{P}[\mathbf{x} = x, \mathbf{i} = i, \mathbf{y} = y, \mathbf{j} = j] = c((x, i), (y, j))$). Recall that the cut s^* cuts an edge $((x, i), (y, j))$ if and only if $x_{s(i)} \neq y_{s(j)}$. It follows that

$$\mathbb{P} \left[\mathbf{x}_{s(i)} = \mathbf{y}_{s(j)} \right] = 1 - \text{val}_{\mathcal{I}^*}(s^*). \quad (48)$$

The claim now follows from Eqs. (47) and (48). \square

Soundness of the reduction from UniqueGames to BalancedSeparator is a reformulation of [KV15, Theorem 11.2] without PCP verifiers:

Lemma 6.10 (Soundness). *(Theorem 11.2 [KV15]) For every $t \in (\frac{1}{2}, 1)$ there exists a constant $b_t > 0$ such that the following holds. Let $\varepsilon > 0$ be sufficiently small and let $\mathcal{I} = \mathcal{I}(w, \pi)$ be an instance of UniqueGames(n, q) and let \mathcal{I}^* be the instance of BalancedSeparator($2^q n, d$) as defined in Section 6.2. If $\text{OPT}(\mathcal{I}) < 2^{-O(1/\varepsilon^2)}$ then $\text{OPT}(\mathcal{I}^*) > b_t \varepsilon^t$.*

We are now ready to prove the main theorem of this section: that no polynomial sized linear program can approximate the BalancedSeparator problem up to a constant factor.

Theorem 6.11. *For every $q \geq 2, \delta > 0, t \in (\frac{1}{2}, 1)$ and $k \geq 1$ there exists a constant $c > 0$ and a demand function $d: E(K_n) \rightarrow \mathbb{R}_{\geq 0}$ for every large enough n , such that $\text{tw}([n]_d) = 2^q - 1$ and*

$$\text{fc}_{\text{LP}} \left(\text{BalancedSeparator}(2^q n, d), \delta + (\log q)^{-1/2}, (\log q)^{-t/2} \right) \geq cn^k.$$

Proof. This statement follows immediately with Lemmas 6.9 and 6.10, together with Theorem 3.2 and Theorem 2.31 with $C_1 = 1 - \delta, S_1 = \frac{1}{q} + \delta, C_2 = \delta + (\log q)^{-1/2}$ and $S_2 = (\log q)^{-t/2}$. Note that the matrices as in Theorem 3.2 are chosen as

$$M_1(\mathcal{I}, s) = \frac{2}{1 - 2\varepsilon}, \quad M_2(\mathcal{I}, s) = \frac{1 + \varepsilon}{1 - 2\varepsilon} \delta$$

with $\varepsilon = (\log q)^{-1/2}$. Since M_1 and M_2 are constant nonnegative matrices $\text{rk}_{\text{LP}} M_1 = \text{rk}_{\text{LP}} M_2 = 1$. \square

Finally, we can prove Theorem 6.2 via choosing the right parameters in Theorem 6.11.

Proof of Theorem 6.2. Straightforward from Theorem 6.11 by choosing $t = \frac{3}{4}, \delta = (\log q)^{-1/2}$ and $q = 2^{(2c_1)^8}$ so that the treewidth of the demand graph is bounded by $c_2 = 2^q - 1 = 2^{2^{(2c_1)^8}} - 1$. \square

7 SDP hardness of MaxCut

We now show that MaxCut cannot be approximated via small SDPs within a factor of $15/16 + \varepsilon$. As approximation guarantees for an instance graph H , we shall use $C(H) = \alpha |E(H)|$ and $S(H) = \beta |E(H)|$ for some constants α and β , and for brevity we will only write α and β .

Theorem 7.1. *For any $\delta, \varepsilon > 0$ there are infinitely many n such that there is a graph G with n vertices and*

$$\text{fc}_{\text{SDP}} \left(\text{MaxCut}(G), \frac{4}{5} - \varepsilon, \frac{3}{4} + \delta \right) = n^{\Omega(\log n / \log \log n)}. \quad (49)$$

Proof. Recall [Sch08, Theorem 4.5] applied to the predicate $P = (x_1 + x_2 + x_3 = 0) \pmod{2}$: For any $\gamma, \delta > 0$, and large enough m , there is an instance \mathcal{I} of Max-3-XOR/0 on m variables with $\text{OPT}(\mathcal{I}) \leq 1/2 + \delta$ but having a Lasserre solution after $\Omega(m^{1-\gamma})$ rounds satisfying all the clauses. By [LRS14, Theorem 6.4], we obtain that for any $\delta, \varepsilon > 0$ for infinitely many m

$$\text{fc}_{\text{SDP}} \left(\text{Max-3-XOR/0}, 1 - \varepsilon, \frac{1}{2} + \delta \right) = m^{\Omega(\log m / \log \log m)}. \quad (50)$$

We reuse the reduction from Max-3-XOR/0 to MaxCut in [TSSW00, Lemma 4.2]. Let x_1, \dots, x_m be the variables for Max-3-XOR/0. For every possible clause $C = (x_i + x_j + x_k = 0)$, we shall use the gadget graph H_C from [TSSW00, Figure 4.1], reproduced in Figure 2. We shall use the graph G , which is the union of all the gadgets $H(C)$ for all possible clauses. The vertices 0 and x_1, \dots, x_m are shared by the gadgets, the other vertices are unique to each gadget.




Figure 2: The gadget H_C for the clause $C = (x_i + x_j + x_k = 0)$ in the reduction from Max-3-XOR/0 to MaxCut. Solid vertices are shared by gadgets, the empty ones are local to the gadget.

A Max-3-XOR/0 instance $\mathcal{I} = \{C_1, \dots, C_l\}$ is mapped to the union $G_{\mathcal{I}} = \bigcup_i H(C_i)$ of the gadgets of the clauses C_i in \mathcal{I} , which is an induced subgraph of G .

A feasible solution, i.e., an assignment $s: \{x_1, \dots, x_m\} \rightarrow \{0, 1\}$ is mapped to a vertex set s^* satisfying the following conditions: (1) $x_i \in s^*$ if and only if $s(x_i) = 1$, (2) $0 \notin s^*$, and (3) on every gadget $H(C)$ the set s^* cuts the maximal number of edges subject to the previous two conditions. It is easy to see that s^* cuts 16 out of the 20 edges of every $H(C)$ if s satisfies C , and it cuts 14 edges if s does not satisfy C . Therefore

$$\text{val}_{G_{\mathcal{I}}}^{\text{MaxCut}(G)}(s^*) = \frac{14 + 2 \text{val}_{\mathcal{I}}^{\text{Max-3-XOR/0}}(s)}{20}, \quad (51)$$

which by rearranging provides the completeness of the reduction:

$$1 - \varepsilon - \text{val}_{\mathcal{I}}^{\text{Max-3-XOR/0}}(s) = 10 \left[\frac{4}{5} - \frac{\varepsilon}{10} - \text{val}_{G_{\mathcal{I}}}^{\text{MaxCut}(G)}(s^*) \right]. \quad (52)$$

It also follows from the construction that $\text{val}_{G_{\mathcal{I}}}^{\text{MaxCut}(G)}$ achieves its maximum on a vertex set of the form s^* : given a vertex set X of G , if $0 \notin X$ then let $s(x_i) = 1$ if $x_i \in X$, and $s(x_i) = 0$ otherwise. If $x_i \in X$ then we do it the other way around: $s(x_i) = 1$ if and only if $x_i \notin X$. This definition makes s^* on the vertices $0, x_1, \dots, x_m$ either agree with X (if $0 \notin X$) or to be complement of X (if $0 \in X$). Then $\text{val}_{G_{\mathcal{I}}}^{\text{MaxCut}(G)}(s^*) \geq \text{val}_{G_{\mathcal{I}}}^{\text{MaxCut}(G)}(X)$ by construction. This means

$$\max \text{val}_{G_{\mathcal{I}}}^{\text{MaxCut}(G)} = \frac{14 + 2 \max \text{val}_{\mathcal{I}}^{\text{Max-3-XOR/0}}}{20}. \quad (53)$$

Thus if $\max \text{val}_{\mathcal{I}}^{\text{Max-3-XOR}/0} \leq 1/2 + \delta$ then $\max \text{val}_{G_{\mathcal{I}}}^{\text{MaxCut}(G)} \leq 3/4 + \delta/10$. Therefore we obtain a reduction with guarantees $C_{\text{MaxCut}(G)} = 4/5 - \varepsilon/10$, $S_{\text{MaxCut}(G)} = 3/4 + \delta/10$, $C_{\text{Max-3-XOR}/0} = 1 - \varepsilon$, $S_{\text{Max-3-XOR}/0} = 1/2 + \delta$, proving

$$\begin{aligned} \text{fc}_{\text{SDP}} \left(\text{MaxCut}(G), \frac{4}{5} - \frac{\varepsilon}{10}, \frac{3}{4} + \frac{\delta}{10} \right) &\geq \text{fc}_{\text{SDP}} \left(\text{Max-3-XOR}/0, 1 - \varepsilon, \frac{1}{2} + \delta \right) \\ &= m^{\Omega(\log m / \log \log m)} = n^{\Omega(\log n / \log \log n)}, \end{aligned} \quad (54)$$

where $n = O(m^3)$ is the number of vertices of G . \square

8 Lasserre relaxation is suboptimal for IndependentSet(G)

Applying reductions within Lasserre hierarchy formulations, we will now derive a new lower bound on the Lasserre integrality gap for the IndependentSet problem, establishing that the Lasserre hierarchy is suboptimal: there exists a linear-sized LP formulation for the IndependentSet problem with approximation guarantee $2\sqrt{n}$, whereas there exists a family of graphs with Lasserre integrality gap $n^{1-\gamma}$ after $\Omega(n^\gamma)$ rounds for arbitrary small γ . While this is expected assuming P vs. NP, our result is unconditional. It also complements previous integrality gaps, like $n/2^{O(\sqrt{\log n \log \log n})}$ for $2^{\Theta(\sqrt{\log n \log \log n})}$ rounds in [Tul09], and others in [AT11], e.g., $\Theta(\sqrt{n})$ rounds of Lasserre are required for deriving the exact optimum.

For IndependentSet(G), the base functions of the Lasserre hierarchy are the indicator functions Y_v that a vertex v is contained in a feasible solution (which is an independent set), i.e., $Y_v(I) := \chi(v \in I)$.

Theorem 8.1. *For any small enough $\gamma > 0$ there are infinitely many n , such that there is a graph G with n vertices with the largest independent set of G having size $\alpha(G) = O(n^\gamma)$ but there is a $\Omega(n^\gamma)$ -round Lasserre solution of size $\Theta(n)$, i.e., the integrality gap is $n^{1-\gamma}$. However $\text{fc}_{\text{LP}}(\text{IndependentSet}(G), 2\sqrt{n}) \leq 3n + 1$.*

Proof. The statement $\text{fc}_{\text{LP}}(\text{IndependentSet}(G), 2\sqrt{n}) \leq 3n + 1$ is [BFPS15, Lemma 5.2]. For the integrality gap construction, we apply Theorem 2.30 with the following choice of parameters. We shall use N for the number of variables, as n will be the number of vertices of G . The parameters q and ε are fixed to arbitrary values. The parameter κ is chosen close to 1, and δ is chosen to be a large constant; the exact values will be determined later. The number of variables N will vary, but will be large enough depending on the parameters already chosen. The parameters β and k are chosen so that the required lower and upper bounds on β are approximately the same:

$$\begin{aligned} k &:= \left\lfloor \frac{(1-\kappa)(\delta-1) \log N - \Theta(\delta \log \log N)}{\log q} \right\rfloor \\ &= \frac{(1-\kappa)(\delta-1) \log N - \Theta(\delta \log \log N)}{\log q} = \Theta(\log N) \end{aligned} \quad (55)$$

$$\beta := \frac{1}{N} \left\lceil \frac{6Nq^k \ln q}{\varepsilon^2} \right\rceil = q^{k+o(1)} = N^{(1-\kappa)(\delta-1) - \Theta(\delta \log \log N / \log N)}. \quad (56)$$

Thus $\beta \geq (6q^k \ln q) / \varepsilon^2$, and for large enough N , we also have

$$\beta \leq N^{(1-\kappa)(\delta-1)} / (10^{8(\delta-1)} k^{2\delta+0.75}).$$

(The role of the term $\Theta(\delta \log \log N)$ in k is ensuring this upper bound. Rounding ensures that k and βN are integers.) By the theorem, there is a k -CSP \mathcal{J} on N variables x_1, \dots, x_N and clauses C_1, \dots, C_m coming

from a predicate P such that $\text{OPT}(\mathcal{I}) = O((1 + \varepsilon)/q^k)$ and there is a pseudoexpectation $\tilde{\mathbb{E}}_{\mathcal{I}}$ of degree at least $\eta N/16$ with $\tilde{\mathbb{E}}_{\mathcal{I}}(\text{val}_{\mathcal{I}}) = 1$. Here

$$m := \beta N = N^{(1-\kappa \pm o(1))(\delta-1)}, \quad (57)$$

$$\eta N/16 = N^{\kappa \pm o(1)}. \quad (58)$$

Let a denote the number of satisfying partial assignments of P . A uniformly random assignment satisfies an a/q^k fraction of the clauses in expectation, therefore $a/q^k \leq \text{OPT}(\mathcal{I}) = O((1 + \varepsilon)/q^k)$, i.e., $a = \Theta(1 + \varepsilon)$.

Let G be the conflict graph of \mathcal{I} , i.e., the vertices of G are pairs (i, s) with $i \in [m]$ and s a satisfying partial assignment s of clause C_i with domain the set of free variables of C_i . Two pairs (i, s) and (j, t) assignments are adjacent as vertices of G if and only if the partial assignments s and t conflict, i.e., $s(x_j) \neq t(x_j)$ for some variable x_j on which both s and t are defined. Thus G has

$$n := am = N^{(1-\kappa)(\delta-1) \pm o(1)} \quad (59)$$

vertices.

Given an assignment $t: \{x_1, \dots, x_N\} \rightarrow [q]$ we define the independent set t^* of G as the set of partial assignments s compatible with t . (Obviously, t^* is really an independent set.) This provides a mapping $*$ from the set of assignments of the x_1, \dots, x_N to the set of independent set of G . Clearly, $\text{val}_G(t^*) = m \text{val}_{\mathcal{I}}(t)$, as t^* contains one vertex per clause satisfied by t . It is easy to see that every independent set I of G is a subset of some t^* , and hence

$$\text{OPT}(G) = m \text{OPT}(\mathcal{I}) = mO((1 + \varepsilon)/q^k) = O(N) = O(n^{1/[(1-\kappa)(\delta-1) \pm o(1)]}). \quad (60)$$

We define a pseudoexpectation $\tilde{\mathbb{E}}_G$ of degree $\eta N/16k$ for G as a composition of $*$ and the pseudoexpectation $\tilde{\mathbb{E}}_{\mathcal{I}}$ of the CSP instance \mathcal{I} :

$$\tilde{\mathbb{E}}_G(F) := \tilde{\mathbb{E}}_{\mathcal{I}}(F \circ *). \quad (61)$$

Recall that $X_{x_j=b}$ is the indicator that b is assigned to the variable x_j , and $Y_{(i,s)}$ is the indicator that (i, s) is part of the independent set. Note that for $s \in V(G)$, we have $Y_{(i,s)} \circ * = \prod_{x_j \in \text{dom } s} X_{x_j=s(x_j)}$ is of degree at most k , and therefore $\deg(F \circ *) \leq k \deg F$, showing that $\tilde{\mathbb{E}}_G$ is well-defined. Clearly $\tilde{\mathbb{E}}_G$ is a pseudo-expectation, as so is $\tilde{\mathbb{E}}_{\mathcal{I}}$.

Now, letting $s \sim C_i$ denote that s is a satisfying partial assignment for C_i :

$$\text{val}_G \circ * = \sum_{(i,s) \in V(G)} Y_{(i,s)} \circ * = \sum_{i \in [m]} \sum_{s \sim C_i} \prod_{x_j \in \text{dom } s} X_{x_j=s(x_j)} = \sum_{i \in [m]} C_i = m \text{val}_{\mathcal{I}}, \quad (62)$$

and hence

$$\tilde{\mathbb{E}}_G(\text{val}_G) = m \cdot \tilde{\mathbb{E}}_{\mathcal{I}}(\text{val}_{\mathcal{I}}) = m = n/a = \Theta(n), \quad (63)$$

showing $\text{SoS}_{\eta N/16k}(G) \geq m$. The number of rounds is

$$\eta N/16k = n^{[\kappa \pm o(1)]/[(1-\kappa)(\delta-1) \pm o(1)]}. \quad (64)$$

From Equations (60), (63) and (64) the theorem follows with an appropriate choice of κ and δ depending on γ . \square

9 From Sherali–Adams reductions to general LP reductions

There are several reductions between Sherali–Adams solutions of problems in the literature. Most of these reductions do not make essential use of the Sherali–Adams hierarchy. The reduction mechanism introduced in Section 3 allows us to directly execute them in the linear programming framework. As an example, we extend the Sherali–Adams reductions from UniqueGames to various kinds of CSPs from [BFPS15] to the general LP case. These CSPs are used in [BFPS15] as intermediate problems for reducing to non-uniform VertexCover and Q -VertexCover, hence composing the reductions here with the ones in [BFPS15] yield direct reductions from UniqueGames to VertexCover and Q -VertexCover.

9.1 Reducing UniqueGames to 1F-CSP

We demonstrate the generalization to LP reductions by transforming the Sherali–Adams reduction from UniqueGames to 1F-CSP in [BFPS15].

Definition 9.1. A *one-free bit* CSP (1F-CSP for short) is a CSP where every clause has exactly two satisfying assignments over its free variables.

Theorem 9.2. With small numbers $\eta, \varepsilon, \delta > 0$ positive integers t, q, Δ as in [BFPS15, Lemma 3.4], we have for any $0 < \zeta < 1$ and n large enough

$$\text{fc}_{\text{LP}}(\text{UniqueGames}_{\Delta}(n, q), 1 - \zeta, \delta) - n\Delta^t q^{t+1} \leq \text{fc}_{\text{LP}}(1\text{F-CSP}, (1 - \varepsilon)(1 - \zeta t), \eta) \quad (65)$$

Proof. Let $V = \{0, 1\} \times [n]$ denote the common set of vertices of all the instances of $\text{UniqueGames}_{\Delta}(n, q)$. The variables of 1F-CSP are chosen to be all the $\langle v, z \rangle$ for $v \in V$ and $z \in \{-1, +1\}^{[q]}$. (Here $\langle v, z \rangle$ stands for the pair of v and z .) Given a $\text{UniqueGames}_{\Delta}(n, q)$ instance (G, w, π) , we define an instance $(G, w, \pi)^*$ of 1F-CSP as follows.

Let v be any vertex of G , and let u_1, \dots, u_t be vertices adjacent to v (allowing the same vertex to appear multiple times). Furthermore, let $x \in \{-1, +1\}^{[q]}$ and let S be a subset of $[q]$ of size $(1 - \varepsilon)q$. We introduce the clause $C(v, u_1, \dots, u_t, x, S)$ as follows, which is an approximate test for the edges $\{v, u_1\}, \dots, \{v, u_t\}$ to be correctly labelled.

$$C(v, u_1, \dots, u_t, x, S) := \exists b \in \{-1, +1\} \forall i \in [t] \forall z \in \{-1, +1\}^{[q]} \begin{cases} \langle u_i, z \rangle = b & \text{if } \pi_{v, u_i}(z) \upharpoonright S = x \upharpoonright S, \\ \langle u_i, z \rangle = -b & \text{if } \pi_{v, u_i}(z) \upharpoonright S = -x \upharpoonright S. \end{cases} \quad (66)$$

We will define a probability distribution on clauses, and the *weight* of a clause will be its probability.

First we define a probability distribution μ_1 on edges of G proportional to the weights. More precisely, we define a distribution on pairs of adjacent vertices (\mathbf{v}, \mathbf{u}) :

$$\mathbb{P}[\{\mathbf{v}, \mathbf{u}\} = \{v, u\}] := \frac{w(v, u)}{\sum_{i,j} w(i, j)}, \quad (67)$$

therefore for the objective of $\text{UniqueGames}_{\Delta}(n, q)$ we obtain

$$\text{val}_{(G, w, \pi)}^{\text{UniqueGames}_{\Delta}(n, q)}(s) = \mathbb{E}[s(\mathbf{v}) = \pi_{\mathbf{v}, \mathbf{u}}(s(\mathbf{u}))] \quad (68)$$

Let μ_1^v denote the marginal of \mathbf{v} in the distribution μ_1 , and $\mu_1^{u|v}$ denote the conditional distribution of \mathbf{u} given $\mathbf{v} = v$.

Now we define a distribution μ_t on vertices $\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t$ such that \mathbf{v} has the marginal distribution μ_1^v , and given $\mathbf{v} = v$, the vertices $\mathbf{u}_1, \dots, \mathbf{u}_t$ are chosen mutually independently, each with the conditional distribution $\mu_1^{u|v}$. Thereby every pair $(\mathbf{v}, \mathbf{u}_i)$ has marginal distribution μ_1 .

Finally, $\mathbf{x} \in \{-1, +1\}^{[q]}$ and $\mathbf{S} \subseteq [q]$ are chosen randomly and independently of each other and the vertices $\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t$, subject to the restriction $|\mathbf{S}| = (1 - \varepsilon)q$ on the size of \mathbf{S} . This finishes the definition of the distribution of clauses, in particular,

$$\text{val}_{(G,w,\pi)^*}^{\text{1F-CSP}}(p) = \mathbb{E} [C(\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t, \mathbf{x}, \mathbf{S})[p]] \quad (69)$$

for all evaluation p .

Feasible solutions are translated via

$$s^*(\langle v, z \rangle) := z_{s(v)}. \quad (70)$$

Soundness of the reduction, i.e., (32-sound) follows from [BFPS15, Lemma 3.4].

Completeness, i.e., (32-complete), easily follows from an extension of the argument in [BFPS15, Lemma 3.5]. The main estimation comes from the fact that the clause $C(v, u_1, \dots, u_t, x, S)$ is satisfied if the edges $\{v, u_1\}, \dots, \{v, u_t\}$ are all correctly labeled and the label $s(v)$ of v lies in S :

$$C(v, u_1, \dots, u_t, x, S)[s^*] \geq \chi[s(v) = \pi_{v, u_i}(s(u_i)), \forall i \in [t]; s(v) \in S]. \quad (71)$$

Let us fix the vertices v, u_1, \dots, u_t and take expectation over x and S :

$$\begin{aligned} \mathbb{E}_{\mathbf{x}, \mathbf{S}} [C(v, u_1, \dots, u_t, \mathbf{x}, \mathbf{S})[s^*]] &\geq (1 - \varepsilon) \chi[s(v) = \pi_{v, u_i}(s(u_i)), \forall i \in [t]] \\ &\geq (1 - \varepsilon) \left(\sum_{i \in [t]} \chi[s(v) = \pi_{v, u_i}(s(u_i))] - t + 1 \right). \end{aligned} \quad (72)$$

We build a nonnegative matrix M out of the difference of the two sides of the inequality. The difference depends only partly on s : namely, only on the values of s on the vertices v, u_1, \dots, u_t . Therefore we also build a smaller variant \tilde{M} of M making this dependence explicit, which will be the key to establish low LP-rank later:

$$\begin{aligned} \tilde{M}_{v, u_1, \dots, u_t}((G, w, \pi), s \upharpoonright \{v, u_1, \dots, u_t\}) &= M_{v, u_1, \dots, u_t}((G, w, \pi), s) \\ &:= \mathbb{E}_{\mathbf{x}, \mathbf{S}} [C(v, u_1, \dots, u_t, \mathbf{x}, \mathbf{S})[s^*]] - (1 - \varepsilon) \left(\sum_{i \in [t]} \chi[s(v) = \pi_{v, u_i}(s(u_i))] - t + 1 \right) \\ &\geq 0. \end{aligned} \quad (73)$$

Taking expectation provides

$$\begin{aligned} \text{val}_{(G,w,\pi)^*}^{\text{1F-CSP}}(s^*) &= \mathbb{E} [C(\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t, \mathbf{x}, \mathbf{S})[s^*]] \\ &= (1 - \varepsilon) \left(\sum_{i \in [t]} \mathbb{P} [s(\mathbf{v}) = \pi_{\mathbf{v}, \mathbf{u}_i}(s(\mathbf{u}_i))] - t + 1 \right) + \mathbb{E} [M_{\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t}((G, w, \pi), s)] \\ &= (1 - \varepsilon) (t \text{val}_{(G,w,\pi)}^{\text{UniqueGames}(n,q)}(s) - t + 1) + \mathbb{E} [M_{\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t}((G, w, \pi), s)], \end{aligned} \quad (74)$$

and hence after rearranging we obtain, no matter what ζ is

$$1 - \zeta - \text{val}_{(G,w,\pi)}^{\text{UniqueGames}(n,q)}(s) = \frac{(1 - \varepsilon)(1 - \zeta t) - \text{val}_{(G,w,\pi)^*}^{\text{1F-CSP}}(s^*) + \mathbb{E} [M_{\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t}((G, w, \pi), s)]}{t(1 - \varepsilon)}. \quad (75)$$

(Note that Equation (75) is not affine due to the last term in the numerator.)

Here the last term in the numerator is the matrix M_2 in the reduction Definition 3.1 (up to the constant factor of the denominator). We show that it has low LP rank:

$$\begin{aligned} \mathbb{E} [M_{\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t}((G, w, \pi), s)] \\ = \sum_{\substack{v, u_1, \dots, u_t \\ f: \{v, u_1, \dots, u_t\} \rightarrow [q]}} \left(\mathbb{P}[\mathbf{v} = v, \mathbf{u}_1 = u_1, \dots, \mathbf{u}_t = u_t] \tilde{M}_{v, u_1, \dots, u_t}((G, w, \pi), f) \right) \\ \cdot \chi(f = s \upharpoonright \{v, u_1, \dots, u_t\}), \end{aligned} \quad (76)$$

i.e., the expectation can be written as the sum of at most $n\Delta^t q^{t+1}$ nonnegative rank-1 factors. Therefore the claim follows from Theorem 3.2. \square

9.2 Reducing UniqueGames to Q - \neq -CSP

Definition 9.3. A *not equal CSP* (Q - \neq -CSP for short) is a CSP with value set \mathbb{Z}_Q , the additive group of integers modulo Q , where every clause has the form $\bigwedge_{i=1}^k x_i \neq a_i$ for some constants a_i .

Theorem 9.4. With small numbers $\eta, \varepsilon, \delta > 0$ positive integers t, q, Δ as in [BFPS15, Lemma 3.4], we have for any $0 < \zeta < 1$ and n large enough

$$\text{fc}_{\text{LP}}(\text{UniqueGames}_\Delta(n, q), 1 - \zeta, \delta) - n\Delta^t q^{t+1} \leq \text{fc}_{\text{LP}}(Q\text{-}\neq\text{-CSP}, (1 - \varepsilon)(1 - 1/q)(1 - \zeta t), \eta) \quad (77)$$

Proof. The proof is similar to that of Theorem 9.2, with the value set $\{-1, +1\}$ consistently replaced with \mathbb{Z}_Q . Let $V = \{0, 1\} \times [n]$ again denote the common set of vertices of all the instances of $\text{UniqueGames}(n, q)$. The variables of Q - \neq -CSP are chosen to be all the $\langle v, z \rangle$ for $v \in V$ and $z \in \mathbb{Z}_Q^{[q]}$.

To simplify the argument, we now introduce additional *hard constraints*, i.e., which have to be satisfied by any assignment. This can be done without loss of generality as these hard constraints can be eliminated by using only one variable from every coset of $\mathbb{Z}_Q \mathbb{1}$ and substituting out the other variables. The resulting CSP will be still a not equal CSP, however this would break the natural symmetry of the structure. Let $\mathbb{1} \in \mathbb{Z}_Q^{[q]}$ denote the element with all coordinates 1. We introduce the *hard constraints*

$$\langle v, z + \lambda \mathbb{1} \rangle = \langle v, z \rangle + \lambda \quad (\lambda \in \mathbb{Z}_Q). \quad (78)$$

Given a $\text{UniqueGames}_\Delta(n, q)$ instance (G, w, π) , we now define an instance $(G, w, \pi)^*$ of Q - \neq -CSP as follows. Let v be any vertex of G , and let u_1, \dots, u_t be vertices adjacent to v (allowing the same vertex to appear multiple times). Furthermore, let $x \in \mathbb{Z}_Q^{[q]}$ and let S be a subset of $[q]$ of size $(1 - \varepsilon)q$. We introduce the clause $C(v, u_1, \dots, u_t, x, S)$ as follows, which is once more an approximate test for the edges $\{v, u_1\}, \dots, \{v, u_t\}$ to be correctly labeled.

$$\begin{aligned} C(v, u_1, \dots, u_t, x, S) &:= \forall i \in [t] \forall z \in \mathbb{Z}_Q^{[q]} \\ &\langle u_i, z \rangle \neq 0 \quad \text{if } \pi_{v, u_i}(z) \upharpoonright S = x \upharpoonright S. \end{aligned} \quad (79)$$

The *weight* of a clause is defined as its probability using the same distribution on vertices $\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t$ as in Theorem 9.2, and randomly and independently chosen $\mathbf{x} \in \mathbb{Z}_Q^{[q]}$ and $S \subseteq [q]$ with $|S| = \varepsilon[q]$. This is the analogue of the distribution in Theorem 9.2, in particular,

$$\text{val}_{(G, w, \pi)}^{\text{UniqueGames}_\Delta(n, q)}(s) = \mathbb{E}[s(\mathbf{v}) = \pi_{\mathbf{v}, \mathbf{u}_i}(s(\mathbf{u}_i))] \quad (i \in [t]), \quad (80)$$

$$\text{val}_{(G, w, \pi)^*}^{1\text{F-CSP}}[p] = \mathbb{E}[C(\mathbf{v}, \mathbf{u}_1, \dots, \mathbf{u}_t, \mathbf{x}, S)[p]]. \quad (81)$$

Feasible solutions are translated via

$$s^*(\langle v, z \rangle) := z_{s(v)}, \quad (82)$$

which clearly satisfy the hard constraints (78).

The reduction is sound by [BFPS15, Lemma 6.9]. For completeness, we follow a similar approach to [BFPS15, Lemma 6.10] and of Theorem 9.2. The starting point is that given a labeling s of (G, w, π) a clause $C(v, u_1, \dots, u_t, x, S)$ is satisfied if the edges $\{v, u_1\}, \dots, \{v, u_t\}$ are correctly labeled, $s(v) \in S$, and $x_{s(v)} \neq 0$:

$$C(v, u_1, \dots, u_t, x, S)[s^*] \geq \chi[s(v) = \pi_{v, u_i}(s(u_i)), \forall i \in [t]; x_{s(v)} \neq 0; s(v) \in S]. \quad (83)$$

Fixing the vertices v, u_1, \dots, u_k and taking expectation over x and S yields:

$$\mathbb{E}_{\mathbf{x}, S} [C(v, u_1, \dots, u_t, \mathbf{x}, S)[s^*]] \geq (1 - \varepsilon)(1 - 1/q)\chi[s(v) = \pi_{v, u_i}(s(u_i)), \forall i \in [t]], \quad (84)$$

where the term $(1 - 1/q)$ arises as the probability of $\mathbf{x}_{s(v)} \neq 0$. The rest of the proof is identical to that of Theorem 9.2, with $1 - \varepsilon$ replaced with $(1 - \varepsilon)(1 - 1/q)$. \square

10 A small uniform LP over graphs with bounded treewidth

Complementing the results from before, we now present a Sherali–Adams like *uniform* LP formulation that solves Matching, IndependentSet, and VertexCover over graphs of bounded treewidth. The linear program has size roughly $O(n^k)$, where n is the number of vertices and k is the upper bound on treewidth. Here uniform means that *the same linear program* is used for all graphs of bounded treewidth with the same number of vertices, in particular, the graph and weighting are encoded solely in the objective function we optimize. This complements recent work [KKT15], which provides a linear program of linear size for a *fixed graph* for weighted versions of problems expressible in monadic second order logic. Our approach is also in some sense complementary to [BM15] where small approximate LP formulations are obtained for problems where the intersection graph of the constraints has bounded treewidth; here the underlying graph of the problem is of bounded treewidth.

Bounded treewidth graphs are of interest, as many NP-hard problems can be solved in polynomial time when restricting to graphs of bounded treewidth. The celebrated Courcelle’s Theorem [Cou90] states that any graph property definable by a monadic second order formula can be decided for bounded treewidth graphs in time linear in the size of the graph (but not necessarily polynomial in the treewidth or the size of the formula).

The usual approach to problems for graphs of bounded treewidth is to use dynamic programming to select and patch together the best partial solutions defined on small portions of the graph. Here we model this in a linear program, with the unique feature that it does not depend on any actual tree decomposition. We call problems *admissible* which have the necessary additional structure, treating partial solutions and restrictions in an abstract way.

Definition 10.1 (Admissible problems). Let n and k be positive integers. Let $\mathcal{P} = (\mathcal{S}, \mathfrak{G}_{n,k}, \text{val})$ be an optimization problem with instances the set $\mathfrak{G}_{n,k}$ of all graphs G with $V(G) \subseteq [n]$ and $\text{tw}(G) \leq k$. The problem \mathcal{P} is *admissible* if

1. *Partial feasible solutions.* There is a set $\mathcal{S} \subseteq \mathcal{P}$ of *partial feasible solutions* and a restriction operation \upharpoonright mapping any partial solution s and a vertex set $X \subseteq [n]$ to a partial solution $s \upharpoonright X$. We assume the identity $(s \upharpoonright X) \upharpoonright Y = s \upharpoonright Y$ for all vertex sets $X, Y \subseteq V(G)$ with $Y \subseteq X$ and partial solutions $s \in \mathcal{S}$. Let $\mathcal{P}_X := \{s \upharpoonright X \mid s \in \mathcal{S}\}$ denote the set of restriction of all feasible solutions to X .

2. *Locality.* The measure $\text{val}_G(s)$ depends only on G and $s \upharpoonright V(G)$ for a graph $G \in \mathfrak{G}_{n,k}$ and a solution $s \in \mathcal{S}$.
3. *Gluing.* For any cover $V(G) = V_1 \cup \dots \cup V_l$ satisfying $E[G] = E[V_1] \cup \dots \cup E[V_l]$ and any feasible solutions $\sigma_1 \in \mathcal{S}_{V_1}, \dots, \sigma_l \in \mathcal{S}_{V_l}$ satisfying

$$\sigma_i \upharpoonright V_i \cap V_j = \sigma_j \upharpoonright V_i \cap V_j \quad \text{for all } i \neq j, \quad (85)$$

there is a unique feasible solution s with $s \upharpoonright V_i = \sigma_i$ for all i .

4. *Decomposition.* Let T be an arbitrary tree decomposition of a graph G with $\text{tw}(G) \leq k$ with bags B_v at nodes $v \in T$. Let $t \in V(T)$ be an arbitrary node of T . Let T_1, \dots, T_m be the components of $T \setminus t$ and $t_i \in V(T_i)$ be the unique node t_i in T connected to t . Clearly, every T_i is a tree decomposition of an induced subgraph $G_i = G[\bigcup_{p \in T_i} B_p]$ of G . Moreover, $B_t \cap V(G_i) = B_t \cap B_{t_i}$.

We require the existence of a (not necessarily nonnegative) function $\text{corr}_{G,T,t}$ such that for all feasible solution s

$$\text{val}_G(s) = \text{corr}_{G,T,t}(s \upharpoonright B_t) + \sum_i \text{val}_{G_i}(s). \quad (86)$$

The decomposition property forms the basis of the mentioned dynamic approach, which together with the gluing property allows the solutions to be built up from the best compatible pieces. The role of the locality property is to ensure that the value function is independent of irrelevant parts of the feasible solutions. In particular, (86) generalizes for the optima, when the restriction σ of the solution to B_t is fixed, this is also the basis of the dynamic programming approach mentioned earlier:

Lemma 10.2. *For any admissible problem \mathcal{P} , with the assumption and notation of the decomposition property we have for any $\sigma \in \mathcal{S}_{B_t}$*

$$\text{OPT}_{s: s \upharpoonright B_t = \sigma}(\text{val}_G(s)) = \text{corr}_{G,T,t}(\sigma) + \sum_i \text{OPT}_{s: s \upharpoonright B_{t_i} \cap B_t = \sigma \upharpoonright B_{t_i} \cap B_t}(\text{val}_{G_i}(s)). \quad (87)$$

Proof. For simplicity, we prove this only for maximization problems, as the proof for minimization problems is similar. By (86), the left-hand side is clearly less than or equal to the right-hand side. To show equality let

$$s_i := \arg \max_{s: s \upharpoonright B_t \cap B_{t_i} = \sigma \upharpoonright B_t \cap B_{t_i}} \text{val}_{G_i}(s)$$

be maximizers. We apply the gluing property for the $s_i \upharpoonright V(G_i)$ and σ .

First we check that the conditions for the property are satisfied. By the properties of a tree decomposition, we have $V(G) = B_t \cup \bigcup_i V(G_i)$ and $E[V(G)] = E[B_t] \cup \bigcup_i E[V(G_i)]$. Moreover, $B_t \cap V(G_i) = B_t \cap B_{t_i}$, and hence $s_i \upharpoonright (B_t \cap V(G_i)) = \sigma \upharpoonright (B_t \cap V(G_i))$. Again by the properties of tree decomposition, for $i \neq j$, it holds $V(G_i) \cap V(G_j) \subseteq B_t$, and hence

$$\begin{aligned} s_i \upharpoonright (V(G_i) \cap V(G_j)) &= s_i \upharpoonright (B_t \cap (V(G_i) \cap V(G_j))) \\ &= \sigma \upharpoonright (B_t \cap V(G_i)) \cap V(G_j) = \sigma \upharpoonright (V(G_i) \cap V(G_j)). \end{aligned} \quad (88)$$

In particular, $s_i \upharpoonright (V(G_i) \cap V(G_j)) = s_j \upharpoonright (V(G_i) \cap V(G_j))$.

Therefore by the gluing property, there is a unique feasible solution s with $s \upharpoonright B_t = \sigma$ and $s \upharpoonright V(G_i) = s_i \upharpoonright V(G_i)$ for all i . Clearly, $\text{val}_G(s)$ is equal to the right-hand side. \square

We are ready to state the main result of this section, the existence of a small linear programming formulation for bounded treewidth graph problems:

Theorem 10.3 (Uniform local LP formulation). *Let $\mathcal{P} = (\mathcal{S}, \mathfrak{G}_{n,k}, \text{val})$ be an admissible optimization problem. Then it has the following linear programming formulation, which does not depend on any tree decomposition of the instance graphs, and has size*

$$\text{fc}_{\text{LP}}(\mathcal{P}) \leq \sum_{X \subseteq V(G), |X| < k} |\mathcal{S}_X|. \quad (89)$$

The guarantees are $C(G) = S(G) = \text{OPT}(G)$. Let V_0 be the real vector space with coordinates indexed by the X, σ for $X \subseteq V(G)$, $\sigma \in \mathcal{S}_X$ with $|X| < k$.

Feasible solutions *A feasible solution $s \in \mathcal{S}$ is represented by the vectors x^s in V_0 with coordinates $x_{X,\sigma}^s := \chi(s \upharpoonright X = \sigma)$.*

Domain *The domain of the linear program is the affine space V spanned by all the x^s .*

Inequalities *The LP has the inequalities $x \geq 0$.*

Instances *An instance G is represented by the unique affine function $w^G: V \rightarrow \mathbb{R}$ satisfying $w^G(x^s) = \text{val}_G(s)$.*

One can eliminate the use of the affine subspace V , by using some coordinates for V as variables for the linear program.

Remark 10.4 (Relation to the Sherali–Adams hierarchy). The linear program above is inspired by the Sherali–Adams hierarchy [SA90] as well as the generalized extended formulations model in [BPZ15]. The LP is the standard $(k-1)$ -round Sherali–Adams hierarchy when \mathcal{P} arises from a CSP: the solution set \mathcal{S} is simply the set of all subsets of $V(G)$, and one chooses $s \upharpoonright X = s \cap X$. The inequalities of the LP are the linearization of the following functions, in exactly the same way as for the Sherali–Adams hierarchy:

$$\chi(s \upharpoonright X = \sigma) := \prod_{i \in \sigma} x_i \prod_{i \in X \setminus \sigma} (1 - x_i).$$

For non-CSPs the local functions take on different meanings that are incompatible with the Sherali–Adams perspective.

With this we are ready to prove the main theorem of this section.

Proof of Theorem 10.3. We shall prove that there is a nonnegative factorization of the slack matrix of \mathcal{P}

$$\tau[\text{OPT}(G) - \text{val}_G(s)] = \sum_{\substack{X \subseteq V(G), |X| < k \\ \sigma \in \mathcal{S}_X}} \alpha_{G,X,\sigma} \cdot \chi(s \upharpoonright X = \sigma), \quad (90)$$

where $\tau = 1$ if \mathcal{P} is a maximization problem, and $\tau = -1$ if it is a minimization problem.

From this, one can define the function w^G as:

$$w^G(x) := \text{OPT}(G) - \tau^{-1} \sum_{\substack{X \subseteq V(G), |X| < k \\ \sigma \in \mathcal{S}_X}} \alpha_{G,X,\sigma} \cdot x_{X,\sigma}, \quad (91)$$

such that it is immediate that w^G is affine, $w^G(x^s) = \text{val}_G(s)$ for all $s \in \mathcal{S}$, and that $\tau[\text{OPT}(G) - w^G(x)] \geq 0$ for all $x \in V$ satisfying the LP inequalities $x \geq 0$. The uniqueness of the w^G follow from V being the affine span of the points x^s , where w^G has a prescribed value.

To show (90), let us use the setup for the decomposition property: Let t be a node of T , and let t_1, \dots, t_m be the neighbors of t , and T_i be the component of $T \setminus t$ containing t_i . Let B_x denote the bag of a node x

of T . Let $G_i := G[\bigcup_{p \in T_i} B_p]$ be the induced subgraph of G for which T_i is a tree decomposition (with bags inherited from T).

We shall inductively define nonnegative numbers $\alpha_{G,X,\sigma,A}$ for $G \in \mathfrak{G}_{n,k}$, $X \subseteq V(G)$, $\sigma \in \mathcal{S}_X$, and $A \subseteq B_t$ satisfying

$$\tau \left[\text{OPT}_{s': s' \upharpoonright A = s \upharpoonright A} (\text{val}_G(s')) - \text{val}_G(s) \right] = \sum_{X \subseteq V(G), \sigma \in \mathcal{S}_X} \alpha_{G,X,\sigma,A} \cdot \chi(s \upharpoonright X = \sigma). \quad (92)$$

This will prove the claimed (90) with the choice $\alpha_{G,X,\sigma} := \alpha_{G,X,\sigma,B_t}$. The help variable A is only for the induction.

To proceed with the induction, we take the difference of Eqs. (86) and (87) with the choice $\sigma := s \upharpoonright B_t$:

$$\tau \left[\text{OPT}_{s': s' \upharpoonright B_t = \sigma} (\text{val}_G(s')) - \text{val}_G(s) \right] = \sum_i \tau \left[\text{OPT}_{s': s' \upharpoonright B_{t_i} \cap B_t = \sigma \upharpoonright B_{t_i} \cap B_t} (\text{val}_{G_i}(s')) - \text{val}_{G_i}(s) \right]. \quad (93)$$

Now we use the induction hypothesis on the G_i with tree decomposition T_i to obtain

$$\tau \left[\text{OPT}_{s': s' \upharpoonright B_t = \sigma} (\text{val}_G(s')) - \text{val}_G(s) \right] = \sum_i \sum_{\substack{X \subseteq V(G_i) \\ |X| < k \\ \sigma \in \mathcal{S}_X}} \alpha_{G_i,X,\sigma,B_{t_i} \cap B_t} \chi(s \upharpoonright X = \sigma). \quad (94)$$

Hence (90) follows with the following choice of the $\alpha_{G,X,\sigma,A}$, which are clearly nonnegative:

$$\alpha_{G,X,\sigma,A} := \begin{cases} \sum_i \alpha_{G_i,X,\sigma,B_{t_i} \cap B_t} & \text{if } X \neq B_t \\ \sum_i \alpha_{G_i,X,\sigma,B_{t_i} \cap B_t} + \tau \left[\text{OPT}_{s': s' \upharpoonright A = \sigma \upharpoonright A} (\text{val}_G(s')) - \text{OPT}_{s': s' \upharpoonright B_t = \sigma} (\text{val}_G(s')) \right] & \text{if } X = B_t. \end{cases} \quad (95) \quad \square$$

We now demonstrate the use of Theorem 10.3.

Example 10.5 (VertexCover, IndependentSet, and CSPs such as e.g., MaxCut, UniqueGames). For the problems MaxCut, IndependentSet, and VertexCover, the set of feasible solutions \mathcal{S} is the set of all subsets of \mathcal{S} . We need no further partial solutions (i.e., $\mathcal{S} := \mathcal{S}$), and we choose the restriction to be simply the intersection

$$s \upharpoonright B := s \cap B.$$

It is easily seen that this makes IndependentSet and VertexCover admissible problems, providing an LP of size $O(n^{k-1})$ for graphs with treewidth at most k . As an example, we check the decomposition property for IndependentSet. Using the same notation as in the decomposition property,

$$\begin{aligned} \text{val}_G(s) - \sum_i \text{val}_{G_i}(s) &= |s| - \sum_i \{ |s \cap V(G_i)| - |E(G_i[s \cap V(G_i)])| \} \\ &= |s \cap B_t| - \sum_i \{ |s \cap B_t \cap V(G_i)| - |E(G_i[s \cap B_t \cap V(G_i)])| \}, \end{aligned} \quad (96)$$

as any vertex $v \notin B_t$ is a vertex of exactly one of the G_i , and similarly for edges with at least one end point not in B_t . Therefore the decomposition property is satisfied with the choice

$$\text{corr}_{G,T,t}(\sigma) := |\sigma \cap B_t| - \sum_i \{ |\sigma \cap B_t \cap V(G_i)| - |E(G_i[\sigma \cap B_t \cap V(G_i)])| \}. \quad (97)$$

For UniqueGames(n, q), the feasible solutions are all functions $[n] \rightarrow [q]$. Partial solutions are functions $X \rightarrow [q]$ defined on some subset $X \subseteq [n]$. Restriction $s \upharpoonright X$ is the usual restriction of s to the subset $\text{dom}(s) \cap X$. This obviously makes MaxCut and UniqueGames(n, q) admissible. The size of the LP is $O(n^{2(k-1)})$ for MaxCut, and $O((qn^2)^{k-1})$ for UniqueGames(n, q).

The Matching problem requires that the restriction operator preserves more local information to ensure that partial solutions are incompatible when they contain a different edge at the same vertex.

Example 10.6 (Matching). The Matching problem has feasible solutions all perfect matchings. The partial solutions are all matchings, not necessarily perfect. The restriction $s \upharpoonright X$ of a matching s to a vertex set X is defined as the set of all edges in s incident to some vertex of X :

$$s \upharpoonright X := \{\{u, v\} \in s \mid u \in B \vee v \in B\}.$$

Now $s \upharpoonright X$ can contain edges with only one end point in X . Again, this makes Matching an admissible problem, providing an LP of size $O(n^k)$ (the number of edges with at most k edges). Here we check the gluing property. Let $V(G) = V_1 \cup \dots \cup V_l$ be a covering (we do not need $E[V(G)] = E[V(G_1)] \cup \dots \cup E[V(G_l)]$), and let σ_i be a (partial) matching covering V_i with every edge in σ_i incident to some vertex in V_i (i.e., $\sigma_i \in \mathcal{S}_{V_i}$) for $i \in [l]$. Let us assume $\sigma_i \upharpoonright V_i \cap V_j = \sigma_j \upharpoonright V_i \cap V_j$, i.e., every vertex $v \in V_i \cap V_j$ is matched to the same vertex by σ_i and σ_j for $i \neq j$. It readily follows that the union $s := \bigcup_i \sigma_i$ is a matching. Actually, it is a perfect matching as $V(G) = V_1 \cup \dots \cup V_l$ ensures that it covers every vertex. Obviously, $s \upharpoonright V_i = \sigma_i$ and s is the unique perfect matching with this property.

Acknowledgements

Research reported in this paper was partially supported by NSF CAREER award CMMI-1452463. Parts of this research was conducted at the CMO-BIRS 2015 workshop *Modern Techniques in Discrete Optimization: Mathematics, Algorithms and Applications* and we would like to thank the organizers for providing a stimulating research environment, as well as Levent Tunçel for helpful discussions on Lasserre relaxations of the IndependentSet problem.

References

- [AAT05] M. Alekhovich, S. Arora, and I. Tourlakis. Towards strong nonapproximability results in the lovász-schrijver hierarchy. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 294–303. ACM, 2005. [13](#)
- [ALN08] S. Arora, J. Lee, and A. Naor. Euclidean distortion and the sparsest cut. *Journal of the American Mathematical Society*, 21(1):1–21, 2008. [19](#)
- [AT11] Y. H. Au and L. Tunçel. Complexity analyses of Bienstock–Zuckerberg and Lasserre relaxations on the matching and stable set polytopes. In *Integer Programming and Combinatorial Optimization*, pages 14–26. Springer, 2011. [2](#), [25](#)
- [AT13] Y. H. Au and L. Tunçel. A comprehensive analysis of polyhedral lift-and-project methods. *arXiv preprint arXiv:1312.5972*, 2013. [2](#)
- [Bar93] F. Barahona. On cuts and matchings in planar graphs. *Mathematical Programming*, 60(1–3):53–68, 1993. [18](#)
- [BCV⁺12] A. Bhaskara, M. Charikar, A. Vijayaraghavan, V. Guruswami, and Y. Zhou. Polynomial integrality gaps for strong SDP relaxations of densest k-subgraph. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 388–405. SIAM, 2012. [13](#)

- [BFPS12] G. Braun, S. Fiorini, S. Pokutta, and D. Steurer. Approximation limits of linear programs (beyond hierarchies). In *53rd IEEE Symp. on Foundations of Computer Science (FOCS 2012)*, pages 480–489, 2012. 2
- [BFPS15] A. Bazzi, S. Fiorini, S. Pokutta, and O. Svensson. No small linear program approximates Vertex Cover within a factor $2 - \epsilon$. *arXiv preprint arXiv:1503.00753*, 2015. 1, 2, 3, 14, 25, 27, 28, 29, 30
- [BM15] D. Bienstock and G. Munoz. LP approximations to mixed-integer polynomial optimization problems. *CoRR, abs/1501.00288*, 2015. 30
- [BOGH⁺06] J. Buresh-Oppenheim, N. Galesi, S. Hoory, A. Magen, and T. Pitassi. Rank bounds and integrality gaps for cutting planes procedures. *Theory Comput.*, 2:65–90, 2006. 13
- [BP15a] G. Braun and S. Pokutta. The matching polytope does not admit fully-polynomial size relaxation schemes. *IEEE Transactions on Information Theory*, 61(10):1–11, 2015. 13, 18
- [BP15b] G. Braun and S. Pokutta. The matching polytope does not admit fully-polynomial size relaxation schemes. *Proceedings of SODA / preprint available at <http://arxiv.org/abs/1403.6710>*, 2015. 18
- [BPZ15] G. Braun, S. Pokutta, and D. Zink. Inapproximability of combinatorial problems via small LPs and SDPs, 2015. 1, 2, 3, 4, 8, 9, 10, 11, 12, 14, 18, 32
- [CKK⁺06] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *computational complexity*, 15(2):94–114, 2006. 19
- [CKN09] J. Cheeger, B. Kleiner, and A. Naor. A $(\log n)^{\Omega(1)}$ integrality gap for the sparsest cut SDP. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pages 555–564. IEEE, 2009. 19
- [CLRS13] S. O. Chan, J. R. Lee, P. Raghavendra, and D. Steurer. Approximate constraint satisfaction requires large LP relaxations. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 350–359. IEEE, 2013. 2, 3, 13
- [CMM09] M. Charikar, K. Makarychev, and Y. Makarychev. Integrality gaps for Sherali–Adams relaxations. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 283–292. ACM, 2009. 2, 13
- [Cou90] B. Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, March 1990. 30
- [FMP⁺12] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf. Linear vs. semidefinite extended formulations: Exponential separation and strong lower bounds. *Proceedings of STOC*, pages 95–106, 2012. 1
- [FMP⁺15] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf. Linear vs. semidefinite extended formulations: Exponential separation and strong lower bounds. *to appear in Journal of the ACM*, 2015. 1
- [Ger91] A. Gerards. Compact systems for T-join and perfect matching polyhedra of graphs with bounded genus. *Operations research letters*, 10(7):377–382, 1991. 18

- [GTW13] A. Gupta, K. Talwar, and D. Witmer. Sparsest cut on bounded treewidth graphs: algorithms and hardness results. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 281–290. ACM, 2013. [14](#), [19](#), [20](#), [21](#)
- [KKT15] P. Kolman, M. Kouteck, and H. R. Tiwary. Extension complexity, mso logic, and treewidth. *arXiv preprint*, July 2015. [2](#), [18](#), [30](#)
- [KV15] S. A. Khot and N. K. Vishnoi. The Unique Games Conjecture, integrality gap for cut problems and embeddability of negative-type metrics into ℓ_1 . *Journal of the ACM (JACM)*, 62(1):8, 2015. [21](#), [23](#)
- [LRS14] J. R. Lee, P. Raghavendra, and D. Steurer. Lower bounds on the size of semidefinite programming relaxations. *arXiv preprint arXiv:1411.6317*, 2014. [1](#), [2](#), [13](#), [24](#)
- [LRST14] J. R. Lee, P. Raghavendra, D. Steurer, and N. Tan. On the power of symmetric LP and SDP relaxations. In *Proceedings of the 2014 IEEE 29th Conference on Computational Complexity*, pages 13–21. IEEE Computer Society, 2014. [12](#)
- [LT03] L. Lipták and L. Tunçel. The stable set problem and the lift-and-project ranks of graphs. *Mathematical programming*, 98(1-3):319–353, 2003. [2](#)
- [Pas12] K. Pashkovich. *Extended Formulations for Combinatorial Polytopes*. PhD thesis, Magdeburg Universität, 2012. [2](#)
- [Rot14] T. Rothvoß. The matching polytope has exponential extension complexity. *Proceedings of STOC*, pages 263–272, 2014. [1](#), [2](#), [13](#)
- [SA90] H. D. Sherali and W. P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM J. Discrete Math.*, 3:411–430, 1990. [32](#)
- [Sch08] G. Schoenebeck. Linear level Lasserre lower bounds for certain k-CSPs. In *Foundations of Computer Science, 2008. FOCS’08. IEEE 49th Annual IEEE Symposium on*, pages 593–602. IEEE, 2008. [2](#), [13](#), [24](#)
- [SK10] R. S. Subhash Khot, Preyas Popat. Approximate Lasserre integrality gap for Unique Games. In *Proc. APPROX/RANDOM*, volume 6302 of *Lecture Notes in Computer Science*, pages 298–311. Springer Berlin Heidelberg, 2010. [13](#)
- [ST99] T. Stephen and L. Tuncel. On a representation of the matching polytope via semidefinite liftings. *Mathematics of Operations Research*, 24(1):1–7, 1999. [2](#)
- [STT07] G. Schoenebeck, L. Trevisan, and M. Tulsiani. Tight integrality gaps for Lovasz-Schrijver LP relaxations of vertex cover and max cut. In *Proc. STOC*, pages 302–310. ACM, 2007. [13](#)
- [TSSW00] L. Trevisan, G. B. Sorkin, M. Sudan, and D. P. Williamson. Gadgets, approximation, and linear programming. *SIAM Journal on Computing*, 29(6):2074–2097, 2000. [24](#)
- [Tul09] M. Tulsiani. CSP gaps and reductions in the Lasserre hierarchy. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 303–312. ACM, 2009. [2](#), [13](#), [25](#)
- [Yan88] M. Yannakakis. Expressing combinatorial optimization problems by linear programs (extended abstract). In *Proceedings of STOC*, pages 223–228, 1988. [1](#)

- [Yan91] M. Yannakakis. Expressing combinatorial optimization problems by linear programs. *J. Comput. System Sci.*, 43(3):441–466, 1991. [1](#)